NBER WORKING PAPER SERIES

HOW MUCH DOES YOUR BOSS MAKE? THE EFFECTS OF SALARY COMPARISONS

Zoë Cullen
Ricardo Perez-Truglia

Working Paper 24841
http://www.nber.org/papers/w24841

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2018

We are thankful for excellent comments from a large number of colleagues, including but not limited to seminar discussants at UCLA-Anderson, Harvard Business School, Wharton, Northwestern, Yale, Caltech, Brown, the Einaudi Institute, Dartmouth, Columbia, Berkeley, Microsoft Research, Boston University, the Luxembourg School of Finance, the Paris School of Economics, the RIDGE Public Economics Workshop, the Barcelona Summer Forum and the AEA Annual Meetings. The collaborating institution provided financial support for the research being conducted. Additionally, Zoe Cullen was a full-time, salaried employee at that institution while the research was being conducted. This project was reviewed and approved by the Institutional Review Board at University of California Los Angeles (IRB#17-001529). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

How Much Does Your Boss Make? The Effects of Salary Comparisons
Zoë Cullen and Ricardo Perez-Truglia
NBER Working Paper No. 24841
July 2018
JEL No. J31,J38,M12,M52,Z13

ABSTRACT

We study how employees learn about the salaries of their peers and managers and how their beliefs about those salaries affect their own behavior. We conducted a field experiment with a sample of 2,060 employees from a multi-billion dollar corporation. We combine rich data from surveys and administrative records with data from the experiment, which provided some employees with accurate information about the salaries of others. First, we document large misperceptions about salaries and identify some of their sources. Second, we find that perceived peer and manager salaries have a significant causal effect on employee behavior. These effects are different for horizontal and vertical comparisons. While higher perceived peer salary decreases effort, output, and retention, higher perceived manager salary has a positive effect on those same outcomes. We provide suggestive evidence for the underlying mechanisms. We conclude by discussing implications for pay inequality and pay transparency.

Zoë Cullen
Rock Center 310
Harvard Business School
Boston, MA 02163
zcullen@hbs.edu

Ricardo Perez-Truglia
Anderson School of Management
University of California, Los Angeles
110 Westwood Plaza
Los Angeles, CA 90095
and NBER
ricardo.truglia@anderson.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w24841

1 Introduction

Employees may take an interest in the salaries of others inside their firm, such as their peers and managers. For example, employees may have concerns about their relative standing, or they may use this information to form expectations about their own future salary. As a result, changing the salary of one employee can affect the behavior of other employees in the firm. These externalities can have important implications for the provision of incentives within the firm and for pay transparency. In this paper, we study salary comparisons using a large-scale, high-stakes field experiment in collaboration with a multi-billion dollar corporation.

The first part of the research design allows us to explore how employees learn about the salaries of others inside the firm. We distinguish between horizontal comparisons (i.e., learning about the salaries of peers) and vertical comparisons (i.e., learning about the salaries of superiors). We use incentivized questions to elicit beliefs about the average salary among one's peers, defined as the employees with the same position title in the same unit. We also elicit beliefs about the average salary among one's managers. To measure salary misperceptions, we compare the perceptions reported in the survey to the actual salaries from the firmâs administrative records. Moreover, our research design explores the sources of misperceptions. For example, to assess whether misperceptions are due to a lack of interest, we use an incentive-compatible method to elicit the willingness to pay for information about peer and manager salaries.

The second part of the research design measures the effects of perceptions about peer and manager salaries on behavior. Our methodology provides causal estimates of the cross-salary elasticities: that is, the elasticities between own behavior and the perceived salaries of other employees. This methodology relies on the combination of three sources of data: a tailored survey, information experiments, and administrative data on the behavior of the employees. For each employee, we flipped a coin to decide whether the employee would receive a piece of information about the average salary among peers. We flipped another coin to decide if the employee would receive a piece of information about the average salary among managers. This information experiment generates exogenous variation in the employeesâ salary perceptions at the end of the survey. We measure whether these exogenous shocks to salary perceptions translate into differences in behavior in the months after survey completion. Moreover, we introduce a simple instrumental-variables model that exploits the experimental variation in beliefs to estimate the cross-salary elasticities.

We conducted the field experiment with a sample of 2,060 employees from a large commercial bank with thousands of employees, millions of customers, and billions of dollars in revenues. This firm is typical in key aspects such as pay transparency and pay inequality.

The close collaboration with the firm, along with its rich technological infrastructure, allowed us to collect unique survey, experimental, and behavioral data. For example, the administrative records include high frequency data on all the swipes in and out of the office, emails sent and received, products sold, and even career outcomes, such as internal transfers, exits, promotions, and raises.

In the first set of results, we show that employees have large and systematic misperceptions about the salaries of their peers and managers. Even though the perceived peer salary is centered around the truth, its mean absolute error is substantial (11.5%). The misperceptions are even higher for manager salary, with a systematic under-estimation of 14.1% and a mean absolute error of 28%. Moreover, we provide suggestive evidence on the sources of these misperceptions. First, we show that most of these misperceptions are not the product of lack of interest: most employees are willing to pay daysâ and even weeksâ worth of salary for a piece of information about the salary of peers or managers. Second, we provide evidence of frictions in information diffusion. When employees are randomly assigned to a signal about the salary of peers or managers, they incorporate such information in their own posterior beliefs. However, the information given to one individual does not travel to other individuals in the employee network, not even to the individualâs closest peer. This finding is also consistent with complementary non-experimental evidence: employees who gossip more and those who are more central in the network do not have more accurate beliefs about the average peer salary. Also consistent with this view, perceptions about average peer salary are not more accurate than what would be expected if employees only knew their own salary.

In the second set of results, we document large effects of salary perceptions on employee behavior. The cross-salary elasticities are statistically and economically significant. For horizontal comparisons, a higher perceived peer salary has negative effects on all of our measures of effort and performance: hours worked, number of emails sent, and sales performance. For example, a 1% increase in perceived peer salary decreases the number of hours worked by 0.94%, implying a behavioral elasticity of -0.94 (p-value=0.04). Vertical comparisons also have a significant effect, but in the opposite direction: a higher perceived manager salary has positive effects on effort and performance. For instance, a 1% increase in perceived manager salary increases the hours worked by 0.15%, implying a behavioral elasticity of 0.15 (p-value=0.04). Moreover, we can confidently reject the null hypothesis that the peer elasticity is equal to the manager elasticity: p-value=0.026 for hours worked, p-value=0.007 for emails sent, and p-value<0.001 for sales performance.

These behavioral elasticities are robust to a number of checks. We find effects of similar magnitude across our three measures of effort and performance. The peer elasticity is -0.943, -0.431, and -0.731, respectively, for hours worked, emails sent, and sales performance, and the manager elasticity is 0.150, 0.130, and 0.106, respectively, for hours worked, emails sent, and sales performance. We show that the results are similar regardless of whether we measure the effects of the information on the behavior at 90 or 180 days post treatment. And we conduct falsification regressions in an event-study fashion: as expected, we find no "effects" of the information treatments on pre-treatment behavior.

We provide suggestive evidence about two traditional mechanisms highlighted in the theoretical and the empirical literature. The first mechanism, salary expectations, posits that employees are motivated by higher perceived peer and manager salaries because they act as a positive signal of the employeesâ own future salary. The second mechanism, social comparisons, posits that employees are discouraged by higher perceived peer and manager salaries because employee morale depends on their relative compensation. To disentangle these mechanisms, we measure the effects of peer and manager salaries on a series of survey outcomes. We find that both peer salary and manager salary have a positive effect on expected future salary. This finding suggests that the salary expectations mechanism is at play in both horizontal and vertical comparisons. We find that perceived peer salary has a large effect on pay satisfaction, job satisfaction, and preferences for within-firm redistribution; in contrast, perceived manager salary has no effect on any of these three outcomes. This evidence suggests that while social comparisons are at play in horizontal comparisons, they do not have a role in vertical comparisons. As complementary evidence, we exploit heterogeneity in the distance between the employee's own position and the managerial position. Consistent with the salary expectations channel, we find that the effects of perceived manager salary are stronger for managerial positions that are closer to the employee's own position. In addition, consistent with the absence of social comparisons, a higher manager salary does not decrease effort or performance even when the managerial position is far above the employee's own position.

Our main finding about the differential effects between horizontal and vertical comparisons has several implications. First, firms may want to refrain from rewarding peers differently, because the reward to one employee has a negative externality on the effort of peers. This externality can help explain why firms provide most of the financial incentives vertically instead of horizontally (Baker et al., 1988). Second, our evidence suggests that female employees may be able to tolerate being paid less than male employees as long as that the male employees hold a different position. This phenomenon may explain why the gender wage gap is large in the vertical margin but small in the horizontal margin (Barth et al., 2017). Third, our findings run counter to the widespread view that social forces tend to compress pay inequality within the firm (Frank, 1984): while this channel may force firms to reduce horizontal inequality, firms do not face resistance to increasing vertical inequality. Fourth, our findings suggest that transparency policies, such as disclosure of CEO pay, may be less effective at curbing inequality than previously thought.

Our paper is related to several strands of literature. While long-standing theoretical literature exists on social comparisons (Frank, 1984; Romer, 1984; Summers, 1988; Lazear, 1989; Akerlof and Yellen, 1990) and salary dynamics (Hirschman and Rothschild, 1973; Lazear and Rosen, 1981; Gibbons and Murphy, 1992; Holmstrom, 1999), the empirical evidence has lagged behind. In a landmark study, Card, Mas, Moretti, and Saez (2012) conducted a field experiment to explore the effects of salary transparency at the University of California. A random sample of employees was sent an email with a link to a public website that listed the salaries of all employees at that university. One week later, the researchers sent another email, this time with a link to an online survey. The authors find that, for workers who had salaries below the peer average, receiving the link to the salary website decreased their job satisfaction and increased their stated intent to find another job.

Other studies have focused on pay transparency and pay inequality using natural experiments (Perez-Truglia, 2015; Rege and Solli, 2015; Dube, Giuliano, and Leonard, 2015; Mas, 2017), laboratory experiments (Charness and Kuhn, 2007; Clark, Masclet, and Villeval, 2010; Gächter and Thöni, 2010; Bracha, Gneezy, and Loewenstein, 2015; Huet-Vaughn, 2017), and field experiments (Cohn, Fehr, Herrmann, and Schneider, 2014; Cullen and Pakzad-Hurson, 2016; Breza, Kaur, and Shamdasani, 2018).[1] For example, Cohn et al. (2014) provide evidence that effort is more affected by cuts in pay relative to peers than by cuts in absolute pay, and that these effects persist over a period of six hours. In addition, Breza et al. (2018) present evidence of aversion to peer inequity: in a month-long experiment with Indian manufacturing workers, the authors estimate that workers give up 9.3% of their earnings to avoid a workplace where they are paid differently than their peers.

Our study makes three contributions to the existing literature. First, while the existing evidence focuses on horizontal salary comparisons, we extend the analysis to include vertical comparisons as well. This extension is important because vertical inequality accounts for the vast majority of the within-firm inequality (Medoff and Abraham, 1980; Baker et al., 1988). For example, in the firm where our experiment was conducted, less than 5% of the inequality in base salaries is horizontal.[2] We find that, since individuals react to horizontal and vertical comparisons in opposite ways, distinguishing between these two types of comparisons is crucial for understanding the causes and consequences of pay inequality. Our

[1] The evidence from laboratory experiments is mixed; for example, while Clark et al. (2010) show that subjects care about the compensation of peers, Charness and Kuhn (2007) find the opposite. Additionally, there is a related literature in psychology (Pritchard et al., 1972; Valenzi and Andrews, 1971; Schmitt and Marwell, 1972).

[2] Details about the inequality decomposition are reported in Appendix A.1.

second contribution is methodological. Rather than estimating the reduced-form effects of transparency, our novel research design identifies cross-salary elasticities by combining an information provision experiment, survey data, and behavioral data. Our third contribution lies in the unique scope of the field experiment – the setting is a high-stakes environment, with thousands of careers and billions of dollars in revenues at stake. Moreover, our close collaboration with the firm allowed us to provide measurements that would have been impossible otherwise, such as the willingness to pay for information and the degree of information diffusion.

This paper is also related to a broader literature on the effects of relative income on outcomes such as happiness (Senik, 2004; Luttmer, 2005) and job satisfaction (Godechot and Senik, 2015; Clark, Frijters, and Shields, 2008). We contribute to this literature by providing novel revealed-preference evidence of relative concerns. This paper is also related to literature showing that individuals have substantial misperceptions of their relative income (Cruces, Perez-Truglia, and Tetaz, 2013; Karadja, Mollerstrom, and Seim, 2017). We contribute to this literature by providing new evidence on the sources of those misperceptions and by showing that such misperceptions can be consequential for economic behavior.

The rest of the paper proceeds as follows. Sections 2 and 3 describe the research and survey designs. Section 4 presents the implementation details and the administrative data. Section 5 presents the results about the formation of salary perceptions. Section 6 presents the results on the effects of salary perceptions on behavior. The last section concludes.

2 Research Design

In this section, we introduce the conceptual framework, hypotheses and the econometric model used to identify the behavioral elasticities.

2.1 Conceptual Framework and Hypotheses

Let subscript i index employees. Let O_i be i's own salary – in the remainder of the paper, we always refer to the base salary (i.e., the salary without any commissions or bonuses). Let Y_i be a form of employee i's behavior, such as the number of hours worked. The most elementary framework would allow the employee's own behavior (Y_i) to be a function of the employee's own salary (O_i):

$$\log(Y_i) = \eta_0 + \eta_{own} \cdot \log(O_i) \tag{1}$$

The parameter η_{own} denotes the own-salary elasticity. For instance, if O_i is the hourly wage and Y_i is the number of hours worked, then η_{own} would be the typical labor supply elasticity. However, an employee's behavior may also depend on the salary of other employees. For instance, an employee may use the salaries of her peers and managers in forming expectations about her own future salary. Also, an employee's morale may depend on the comparison of her own salary to that of peers and managers.

In a fully saturated model, we would allow the right-hand side of equation (1) to depend on the salaries of every other employee in the corporation. Since the estimation of such a model would be unfeasible, we focus on a stylized model that allows behavior to depend on two additional terms: the average salary among peers (i.e., employees who share the same position title and organizational unit) and the average salary among managers (i.e., employees who hold positions at a higher level, supervising the employee, and whose position the employee could aspire to be promoted to). For instance, the peers of a teller are the other tellers working in the same branch, and one managerial position is the teller supervisor position. For a junior researcher at the investment bank unit, her peers are the other junior researchers, and one managerial position is the senior researcher position.

Let P_i be i's perception about the average salary of i's peers. Let M_i be i's perception of the average salary among i's managers. The following equation establishes the potential relationship between behavior and salaries:

$$\log(Y_i) = \eta_0 + \eta_{own} \cdot \log(O_i) + \eta_{peer} \cdot \log(P_i) + \eta_{mgr} \cdot \log(M_i) \tag{2}$$

In addition to the own-salary elasticity (η_{peer}), this framework allows employees to react to their perceptions of the salaries of other employees through η_{peer} and η_{mgr}. We denote these two parameters as the cross-salary elasticities.

Cross-salary elasticities may be significant for multiple reasons. On the one hand, *career concerns* posit that employees use their perceptions about the salaries of others to update beliefs and expectations about their own careers. For example, employees may use the salaries of peers and managers to form expectations about their own future salary (Lazear and Rosen, 1981; Gibbons and Murphy, 1992; Holmstrom, 1999). According to the "tunnel effect" (Hirschman and Rothschild, 1973), the expected own future salary should increase with perceived peer and manager salaries. As a result, this mechanism would predict positive cross-salary elasticities ($\eta_{peer} > 0$ and $\eta_{mgr} > 0$). On the other hand, *social concerns* posit that individuals care directly about the salaries of others. For example, higher perceived salaries of peers may demoralize employees (Frank, 1984; Romer, 1984; Summers, 1988; Akerlof and Yellen, 1990; Lazear, 1989; Card et al., 2012). As a result, this mechanism predicts negative cross-salary elasticities ($\eta_{peer} < 0$ and $\eta_{mgr} < 0$).

In the results section we discuss the potential mechanisms in more detail and provide suggestive evidence about them. The results section also discusses and relaxes two implicit functional form assumptions from the stylized model: (log-log) linearity and symmetry.

2.2 Econometric Model

In this section, we present the empirical framework used for identification.

Obtaining causal estimates of the cross-salary elasticities η_{peer} and η_{mgr} is challenging, because a simple regression of behavior on perceived salaries would be subject to the usual concerns about omitted variable biases. For instance, individuals with lower perceived peer salary may have lower ability, which would create a spurious elasticity $\eta_{peer} < 0$. Our empirical framework exploits the random variation in beliefs induced by the information provision experiments to estimate the cross-salary elasticities.

To understand the intuition behind this model, consider a pair of employees who have the same bias about perceived peer salary: both of them underestimate the actual peer salary by 10%. Then, we randomly assign information about the true peer salary to one of these two employees. The first stage of the regression measures the effect of the information provision on beliefs. Assume that, relative to the individual who does not get the information, the individual who receives the information ends up with perceived peer salary that is (say) 5% higher. The reduced form regression follows the behavior of these two individuals after they complete the survey. Assume that, relative to the individual who does not get the information, the individual who receives the information ends up exerting 2.5% lower effort. The instrumental variables model infers the elasticity by putting these two results together. In this example, the information provision increases perceived peer salary by 5% and reduces effort by 2.5%, which implies a cross-salary elasticity of -0.5 ($= \frac{-2.5}{5}$).

The above example centered on a pair of individuals who underestimated their peer salary by 10%. The following instrumental variables model generalizes this logic. Let P_i^{prior} denote the prior belief about the average salary of peers – that is, the belief right before the individual reaches the information-provision experiment. Let P_i^{signal} be the value of the signal on average salary that is randomly assigned in the information-provision experiment. Let T_i^P be a dummy variable that takes the value 1 if the individual is shown the signal and 0 if not. Denote P_i^{post} as the corresponding posterior belief – that is, the perceived peer salary right after the information-provision experiment was conducted.

The first stage of the instrumental variables model corresponds to belief updating. When priors and signals are normally distributed, Bayesian learning implies that the mean of the posterior belief should be a weighted average between the signal and the mean of the prior belief. In the case of perceived peer salary, Bayesian learning can be summarized as follows:

$$P_i^{post} = \alpha \cdot \left(P_i^{signal} - P_i^{prior}\right) + P_i^{prior} \tag{3}$$

This simple Bayesian model or slight variations of it have been found to fit the data well in several contexts, such as inflation expectations (Armantier et al., 2016; Cavallo et al., 2017) and perceptions of relative income (Bottan and Perez-Truglia, 2017). We can introduce the information provision experiment to this learning specification:

$$P_i^{post} = \alpha \cdot \left(P_i^{signal} - P_i^{prior}\right) \cdot T_i^P + \beta \cdot \left(P_i^{signal} - P_i^{prior}\right) + P_i^{prior} \tag{4}$$

The parameter α represents the learning rate, which ranges from 0 (individuals ignore the signal) to 1 (individuals fully adjust to the signal). In turn, the parameter β controls for any spurious reversion to the signal.[3] Moreover, let $M_i^{prior}, M_i^{signal}, T_i^M, M_i^{post}$ be defined as $P_i^{prior}, P_i^{signal}, T_i^P, P_i^{post}$, but for perceived manager salary instead of perceived peer salary. Then, we can apply the identical Bayesian learning model to the beliefs about manager salary.

Equation (4) captures the effects of the information provision experiment on the posterior beliefs. The instrumental variables regression simply exploits that exogenous variation in beliefs for the identification of causal effects. Let Y_i^{post} denote some average behavior in the period beginning with the information provision experiment and ending some time later. For instance, Y_i^{post} could be the average number of hours spent in the office in the 90 days after the information provision. The instrumental-variables model is the following:

$$\begin{aligned} \log\left(Y_i^{post}\right) &= \pi_0 + \eta_{peer} \cdot \hat{P}_i^{post} + \eta_{mgr} \cdot \hat{M}_i^{post} + \\ &+ \pi_1 \cdot \left(P_i^{signal} - P_i^{prior}\right) + \pi_2 \cdot \left(M_i^{signal} - M_i^{prior}\right) + \pi_3 \cdot P_i^{prior} + \pi_4 \cdot M_i^{prior} + \epsilon_i \end{aligned} \tag{5}$$

$$\begin{aligned} P_i^{post} &= \nu_0 + \nu_1 \cdot \left(P_i^{signal} - P_i^{prior}\right) \cdot T_i^P + \nu_2 \cdot \left(M_i^{signal} - M_i^{prior}\right) \cdot T_i^M + \\ &+ \nu_3 \cdot \left(P_i^{signal} - P_i^{prior}\right) + \nu_4 \cdot \left(M_i^{signal} - M_i^{prior}\right) + \nu_5 \cdot P_i^{prior} + \nu_6 \cdot M_i^{prior} + \xi_i^1 \end{aligned} \tag{6}$$

$$\begin{aligned} M_i^{post} &= \mu_0 + \mu_1 \cdot \left(P_i^{signal} - P_i^{prior}\right) \cdot T_i^P + \mu_2 \cdot \left(M_i^{signal} - M_i^{prior}\right) \cdot T_i^M + \\ &+ \mu_3 \cdot \left(P_i^{signal} - P_i^{prior}\right) + \mu_4 \cdot \left(M_i^{signal} - M_i^{prior}\right) + \mu_5 \cdot P_i^{prior} + \mu_6 \cdot M_i^{prior} + \xi_i^2 \end{aligned} \tag{7}$$

The relevant exclusion restrictions for the identification of η_{peer} and η_{mgr} are

[3] For example, if individuals take some additional time to think when asked a question a second time, they tend to revert to the truth, even if they did not actually receive any further information from the experimenter.

$$E\left[\left(P_i^{signal} - P_i^{prior}\right) \cdot T_i^P \cdot \epsilon_i\right] = 0 \text{ and } E\left[\left(M_i^{signal} - M_i^{prior}\right) \cdot T_i^M \cdot \epsilon_i\right] = 0 \quad (8)$$

The random assignment of $\{T_i^P, T_i^M\}$ satisfies these exclusion restrictions. In practice, we include a set of additional control variables to reduce the variance of the error term and thus improve the precision of the estimates: own income (in logs), tenure (in logs), dummies for performance evaluations in the previous year, and, following the standard practice in field experiments (McKenzie, 2012), the pre-treatment outcomes.

Last, we can exploit the timing of the intervention to provide a falsification test in an event-study fashion. Let Y_i^{prior} denote the average behavior in the period prior to the information-provision experiment (rather than the period after it, as in Y_i^{post}). We can estimate the same instrumental variables model from above, but using Y_i^{prior} instead of Y_i^{post} as dependent variable. Intuitively, the information-provision experiment should not affect the behavior in the pre-treatment period, because the individuals have not yet been exposed to the information. Thus, we expect the cross-salary elasticities to be zero in this falsification regression.

3 Survey Design

A sample of the full online survey is included in Appendix C – to protect the identity of the firm, we removed identifying information, including formatting. In this section, we discuss the most important aspects of the survey design.

3.1 Training

The first survey module was designed to provide the respondents with an understanding of the rest of the survey. It begins with an explanation of how the incentivized questions work and why responding honestly is in the respondent's best interest. To cement this knowledge, we included some practice questions on topics that are unrelated to salaries.

This module also provides the definition of salary used in the rest of the survey. We focus on monthly base salary, that is, the salary before any additions or deductions, such as taxes, allowances, commissions, or bonuses. According to interviews with administrators from the Human Resources department and employees who were not participating in the experiment, base salary is the part of the employee compensation that is most salient for them. For example, when a new employee joins the bank, the monthly base salary is the most important figure written in the contract. Moreover, the base salary accounts for nearly all of the total compensation for the subjects in our sample.[4]

To confirm that respondents understood the definition of salary, we asked them to estimate their own salary for the month of March 2017. Subjects were told that they would receive a reward for accuracy. All the accuracy rewards in the survey were set up using the traditional quadratic loss function calibrated to award up to $2.61 per question (in 2017 USD PPP). In the following screen we showed their guess and the truth side-by-side. If the respondent's guess was not within 5% of the truth, we showed them an additional screen explaining the definition again.[5]

3.2 Salary Perceptions

We start with the module on perceived peer average, which followed the structure below:

- **Step 1 (Elicit Prior Belief):** We asked respondents about the average monthly base salary among their peers. The survey instructions are explicit about the group definitions: we stated the full position title, the full name of the unit, and the number of employees currently working in that peer group. To elicit truthful responses, we offered a reward for accuracy. We did not inform subjects whether they got any of the specific questions right or wrong to prevent them from inferring any information about the average salaries from the rewards. Furthermore, the rewards were paid jointly with other participation fees that were randomly determined, so that the total payment received was a weak signal of the overall accuracy of the responses during the survey. After the subject provided the mean belief about the average salary, we elicited the probability beliefs over a series of bins around the respondent's guess – this question was also incentivized.

- **Step 2 (Elicit Willingness to Pay):** We offered respondents the opportunity to buy the following piece of information: the average salary over a random sample of five peers. To elicit this information in an incentive-compatible way, we employed the multiple price list variation of Becker-DeGroot-Marschak (Andersen et al., 2006). This method consists of having respondents make choices in five hypothetical scenarios. In each scenario, the respondent can choose to either see the piece of information or add

[4] Because of the sensitive nature of the data, we refrain from providing exact estimates. After base salary, the second source of compensation for individuals who have some form of sales role is sales commissions. Other forms of performance pay can be substantial for employees at the highest levels of the firm, but those employees were excluded from participating in the survey.

[5] This question was also intended to convey that the surveyor already knows the salary of the respondent, thus preventing any inclination to misreport salaries in an effort to avoid revealing one's own salary to the surveyor.

a certain amount to their survey rewards (i.e., the "price" of the information). Since all employees are also clients of the bank, these rewards amounts could be deposited directly into respondents' current account. The five scenarios differ in the information price: \$1.3, \$6.5, \$26.1, \$130.5, and \$652.3 (in 2017 USD PPP).[6] We explained to subjects that choosing truthfully was in their best interest because there was a 1% probability that one of the five scenarios would be randomly selected to be executed. For the 2% of respondents who had their scenarios executed (1% for the peer salary and 1% for the manager salary), the survey was automatically terminated; thus, they are excluded from the subject pool. The other 98% of respondents continued with the next step.

- **Step 3 (Information-Provision Experiment):** With a 50% probability, subjects were shown the average salary among a random sample of five peers. To avoid individuals making inferences from the act of receiving information, we made the randomization explicit. In a first screen, we let the respondents know that a group of individuals participating in this survey would be randomly chosen to receive some information about the average salary of peers for free. In the following screen, we let the subjects know whether they were chosen to receive the information or not.

- **Step 4 (Elicit Posterior Belief):** We gave the subjects the opportunity to revise their guess about the average salary among peers. To avoid subjects making inferences based on the opportunity to re-elicit their guesses, we explicitly noted that all survey participants automatically had this opportunity, regardless of their initial guesses.

After the module on peer perceptions was over, subjects were shown a module on manager perceptions – that is, this module had the same four steps as the previous module, except it was about the average manager salary instead of the average peer salary. For this step, we assigned each employee to a specific managerial position using multiple sources of administrative data and based on the following criteria: i. The managerial position must be higher than the respondent's position; ii. There must be at least one manager in that position in the respondent's unit; iii. That manager must have oversight over the respondent, such as conducting performance evaluations or approving leaves of absence.[7] The distance between the respondent and the managerial position can be important for understanding the mechanisms at play. For this reason, before eliciting the perceived manager salary, we included two questions to assess the employee's perceived distance to the manager: the number

[6] We calibrated this scale using a small pilot survey that elicited willingness to pay for information with an open-ended and non-incentivized question.

[7] For some employees, we had to pick one manager from among multiple managerial positions that satisfied all these criteria.

of promotions needed to reach that position, and the likelihood of being promoted to that position within five years.

With respect to the information-provision experiment, we cross-randomized the two pieces of information, which resulted in four treatment groups that were equally likely: one group received a signal about the average salary of peers but no salary information about their manager; one group received a signal about the salary of their manager but not peers; one group received information about both their peers' and manager's salary; and one group received no salary information.

We focused on these two groups of individuals, peers and managers, based on interviews with employees and with managers from the Human Resources division. To validate this design choice, at the end of the survey we asked employees to rank how interested they would be in learning about the average salary in different positions: their same position, the positions right above their own level, two or more levels above, and other positions. The data confirm that employees were mostly interested in the average salaries among peers and managers: roughly 50% of subjects ranked their own position first, 45% of subjects ranked higher positions first, and less than 5% of respondents ranked other positions first.[8]

3.3 Survey Outcomes

The main focus of the paper consists of measuring the effects of salary perceptions on behavior. Additionally, to provide some complementary evidence on the potential mechanisms at play, we included a few questions after the perception module to be used as survey outcomes.

The first three questions provide suggestive evidence of social concerns. First, we follow a literature that uses self-reported employee satisfaction to proxy for employee morale (Clark and Oswald, 1996; Card et al., 2012). We included one question about *pay satisfaction*: "How satisfied are you with your current salary?" Responses to this question used a 5-point scale from very dissatisfied (1) to very satisfied (5). Additionally, we included a question on overall *job satisfaction*: "Taking all the aspects of your job into account, how satisfied are you with your current job?" This question used the same 5-point response scale. The third question corresponded to an adaptation of a question that is widely used in the literature on preferences for redistribution (Perez-Truglia, 2015). *Redistribution Preferences* elicited preferences for within-firm redistribution, rather than the usual within-country redistribution: "Across the thousands of [Bank Name] employees, salaries vary with the nature of work, education,

[8] Another design choice was to focus on group averages, as opposed to other characteristics of the within-group distribution, such as the median, minimum, or maximum. This design choice was also based on interviews with employees and Human Resources managers. If anything, to the extent that our choice of specification missed other important characteristics of the salary distribution, our baseline model would underestimate the effects of salary comparisons.

experience, responsibilities, etc. What do you think of wage differentials in the company today?" The possible answers were (1) they are too small, (2) they are adequate, and (3) they are too large. As a result, higher values of this outcome indicate higher preferences for redistribution inside the firm.

The second set of three questions provided suggestive evidence about some career concerns. The first two questions were intended to assess whether individuals use the salaries of other employees to form expectations about their own future salary. For that, we elicited each employee's *expected future salary* 1 year and 5 years ahead. We could not incentivize these questions comparing the responses to the actual future salaries because we would need 1 and 5 years to determine the ex post accuracy. Instead, we incentivized these two questions by comparing employeesââ guesses to our own predictions based on historical data. The third question was intended to assess whether individuals use the salaries of other employees to infer their relative productivity. At the end of every year, each employee is given a productivity rating on a scale from D to A+. We elicited the individual's perception about the share of employees for each rating during the last yearly review. We incentivized this question by rewarding individuals for accurate responses. With these perceived shares and the employee's own rating, we can infer the employee's *perceived productivity rank*.

4 Institutional Context, Data, and Subject Pool

4.1 Institutional Context

We conducted the experiment in collaboration with a private commercial bank in Asia. To keep the identity of the bank secret, we refrain from providing exact information about its characteristics. This bank has millions of customers, billions of dollars in assets and in revenues, and thousands of employees. These employees are based in two headquarter offices and hundreds of branches around the country.

This bank is comparable to large firms around the world in two key respects: degree of pay inequality and degree of pay transparency.

Regarding pay inequality, the ratio between the 10th and 90th percentile of base salary is 0.21. In comparison, the ratio is 0.19 for the average medium-sized firm in the United States (Bloom et al., 2015). We can also benchmark the degree of inequality by looking at salary differential across specific positions. In the firm where the experiment was conducted, the salary between the salary of a senior relationship manager and its subordinate, a personal retail banker, is around 1.5. According to 2017 data from Glassdoor, the corresponding ratio for Bank of America is also 1.5. The inequality in this firm is also typical in that only a minority of it is horizontal (Medoff and Abraham, 1980; Baker et al., 1988): only 5% of the inequality in base salaries is horizontal.[9] And according to anecdotal evidence, these horizontal inequalities seem to follow the typical mix of meritocratic and non-meritocratic reasons.[10]

Regarding pay transparency, the bank does not disclose information about the salaries of other employees, and discourages employees from discussing their salaries with coworkers. These policies are typical in firms around the world; for example, in a survey of private sector employees from the United States, more than 60% report that their employer discourages or prohibits employees from discussing salaries with coworkers (Hegewisch et al., 2011).[11] A complementary measure of pay secrecy is given by the reported frequency with which employees discuss salaries with coworkers. According to a 2017 survey of Americans 18-36 years old, around 30% of them discuss their salaries with coworkers.[12] According to a different but related question from our own survey, 55% of employees report discussing salaries with coworkers.[13]

4.2 Survey Implementation

We started with the universe of thousands of employees.[14] We were asked to focus on the two main units of the company. Some of these employees were disqualified from participating in the subject pool. For instance, if an employee had a position title that is unique in the firm, she would not have any peers according to our definition. Similarly, for some employees, it was difficult to identify a managerial position. After accounting for employees who were ineligible to participate, we were left with 3,841 individuals who could be invited to take the survey.

A sample of the invitation email (stripped of formatting and identifying information)

[9]Details for the inequality analysis are reported in Appendix A.1.

[10]For example, employees may get raises after good performance reviews. While some employees may qualify for these performance raises because of their effort or productivity, other employees may qualify because they flatter their managers. Similarly, employees who attract job offers from other firms may receive raises as part of retention packages. While some employees may attract outside offers due to their innate talent, others may be just good at gaming the system.

[11]Like in other firms, the company where we conducted the experiment discloses some information about pay but this information is too vague to form a decent guess for the average salaries of peers and managers. For instance, the firm discloses the existence of a 10-point payband scale, but the minimum and maximum salaries across these paybands are overlapping and kept secret.

[12]Source: "Ask Me How Much Money I Make," The Wall Street Journal, Oct. 26, 2017. Note that, to the extent that it is frowned upon by the employer, this type of behavior is probably under-reported in surveys.

[13]More precisely, 45% of employees reported never talking about salaries; 16%, once a year; 31%, a few times a year; 6%, once a month; and the remaining 2%, once a week or more often.

[14]This universe excludes recent hires – that is, employees who had been hired less than 6 months prior to the time of the data collection.

is presented in Appendix A.2. The invitation email did not provide any details about the content of the survey. The invitation show endorsement from three executives from the bank, and advertised survey participants would receive, on average, $30 as participation rewards. The survey was not compulsory, but employees were encouraged to participate â indeed, the unit heads reached out to their employees by email to encourage participation in our survey.

The email invitations were sent gradually over time. This staggered timing was designed with two goals in mind: first, we wanted to smooth the burden of the survey over time; second, surveying subjects at different points in time allows us to study information diffusion. in the survey responses span from the first week of April 2017 to the first week of June 2017. Of the 3,841 invitations sent out, 2,060 individuals completed the main module of the survey, corresponding to a 53.6% response rate.[15]

4.3 Descriptive Statistics and Randomization Balance

Table 1 presents some descriptive statistics about the subject pool. Column (1) corresponds to the entire sample of 2,060 survey respondents: 73% of them are female, 86% finished college or a higher degree, and on average they are 29 years old and have been working at the firm for the last 5 years. The last rows of Table 1 report the averages for the own, peer, and manager salaries – due to the sensitive nature of this data, we use an arbitrary unit of measurement that we do not disclose.

The salary data indicate that the degree of horizontal inequality is small: the mean absolute difference between one's own salary and the average peer salary is 11.7% of the own salary. In contrast, the degree of vertical inequality is large: the mean absolute difference between one's own salary and the average manager salary is 315% of the own salary. We chose managerial positions that are relatively close to the employee's own position: the average employee expects to need 3.65 promotions to reach the managerial position and thinks there is a probability of 60.4% of being promoted to that position within the following five years. If we had chosen managerial positions that are more distant, then the degree of vertical inequality would be even higher.

Subjects were cross-randomized to receive information about peers and managers, which resulted in four treatment groups. In columns (2) through (5) of Table 1, we break down the descriptive statistics within each of the four treatment groups. The last column reports

[15]This sample already excluded the individuals who were randomly assigned to have their choices in the information-shopping scenarios executed, for whom the survey was automatically terminated. This final sample also excluded 15 subjects with the most extreme prior beliefs (most likely due to typos). Last, there was a small attrition (less than 3%) between the information provision screen and the corresponding posterior beliefs. We find that this attrition is orthogonal to the instrumental variables used for the experimental analysis.

p-values for the null hypothesis that the average characteristics are the same across all four treatment groups. The results show that, consistent with successful random assignment, the observable characteristics are balanced across the four treatment groups.

The subject pool spans employees from the lowest pay bands, such as tellers, to some of the highest pay bands, such as unit directors. The pool includes employees with different roles such as tellers, sellers, clerks, and receptionists. The subject pool includes employees from the two headquarter offices and from hundreds of branches located all over the country. In Appendix A.2, we provide more details about the subject pool – for instance, we show that the subject pool is representative of the universe of employees in many observable characteristics.

4.4 Behavioral Outcomes

We collaborated with the different units of the organization to create a centralized and anonymous database covering many forms of employee behavior. The main outcomes are the effort and performance of employees. We have two measures of effort. First, for employees working in the headquarter offices (29% of the sample), they must clock-in and -out from the office using an electronic card-swipe system, which is strictly enforced by security personnel. We use these time stamps to calculate the *hours in the office* on a daily basis. The second measure of effort is based on email data. We collect real-time data on the emails sent and received by all employees, which we use to create the variable *number of emails sent* on a daily basis. The advantage of this measure over the alternative of hours worked is that it is available for 100% of the employees. While the number of emails may not be a good measure of effort in other contexts, it seems to be a good proxy in our context and possibly even better than the numbers of hours worked. For security reasons, employees can only access their work email account from their office computers, implying that they can only send emails while at the office. Employees are strongly discouraged from using their work email account for matters unrelated to work. Last but not least, employees need to send emails to clients or coworkers for most of their duties, such as contacting clients, obtaining information about clients, or getting approval for a new loan, credit card or mortgage. Consistent with this suggestive evidence, the number of emails is positively correlated to our alternative measure of effort, the number of hours spent in the office (correlation coefficient of 0.24, p-value<0.001).

Measuring performance in an objective and standardized manner is quite challenging for many positions within the bank.[16] However, for the 38% of employees who have a sales role, we can measure performance based on their sales revenues. The bank uses a sophis-

[16]For instance, a Human Resources employee is involved in various tasks such as identifying new hires, processing paperwork, and dealing with complaints from existing employees. The performance at each of those tasks is difficult to measure with objective data.

ticated formula to aggregate an employee's sales over all products (e.g., credit cards, loans, mortgages). We use this data to construct a *sales performance index* on a monthly basis.

In addition to effort and sales performance, information about peer and manager salaries may affect some career outcomes – for instance, an employee may react by quitting her job or renegotiating her salary. Using multiple sources of administrative data, we constructed four main career outcomes. The dummy variable *quit* takes the value 1 if the employee leaves the firm. The variable *transfer* takes the value 1 if the employee transfers to another unit inside the firm. The variable *salary* is equal to the salary of the employee at the end of the corresponding period. Similarly, the variable *changed title* takes the value 1 if the individualâs position title changed.[17]

The baseline specification defines the post-treatment outcomes as the average outcome in the period starting from the survey date and ending 90 days later – in the results section, we consider alternative time windows.[18] We began collecting data on these behavioral outcomes three months before launching the survey. As a result, in addition to post-treatment outcomes, we can compute the corresponding pre-treatment outcomes, which can be used as control variables to improve precision as well as for falsification tests.

5 Results: Formation of Perceptions about Peer and Manager Salary

In this section, we document the accuracy of perceptions about peer and manager salaries and provide evidence about the potential sources of misperceptions.

5.1 Accuracy of Prior Beliefs

Figure 1 shows the basic evidence about misperceptions of peer and manager salaries. As a benchmark, Figure 1.a shows the misperceptions about own salary, that is, the difference between the individual's perceived own salary and actual own salary. There are nearly no misperceptions about own salary: around 80% of respondents report an own salary that is

[17] The variables *quit* and *transfer* are based on daily data, while the variables *salary* and *changed title* are based on monthly data.

[18] This definition applies to variables constructed with daily data. For variables constructed with monthly data, the post-treatment period corresponds to the month of the survey and the following two months. This specification can lead to an attenuation bias because individuals who respond to the survey on the first day of the month (who were exposed to the information for a full month) would be coded the same as individuals responding on the last day of the month (who were exposed for one day). Last, for the small fraction of employees that leave the company during the relevant time window, we use the average outcome between the survey date and the exit date.

within 5% of the true salary. This outcome confirms our prior belief that this definition of salary is salient. The remaining 20% of employees probably misunderstood the definition of base salary. To make sure that these misunderstandings do not extend to the rest of the survey, these employees were retrained about the definition of base salary. After this stage, we asked these employees if they agreed with our measure of own salary, and 87% of them responded affirmatively.

Figure 1.b shows misperceptions about average peer salary. Relative to own salary, these subjects have accurate perceptions about peer salary much less commonly: only 32% of perceptions are within 5% of the truth. The mean absolute difference between the perceived average and the actual average (i.e., the mean absolute error) is 11.5%. These misperceptions are economically significant, but they are not slanted: approximately as many people over-estimate the average peer salary as the number of people who under-estimate it, resulting in an average over-estimation of peer salary of just 2.5% (p-value<0.01).

Figure 1.c shows misperceptions about average manager salary. Only a small share (12%) of respondents guess the average manager salary within 5% of the truth. The mean absolute error for perceived manager salary (28%) is substantially higher than that of peer salary (11.5%). Also, while there was not a systematic bias in peer salary, the average employee under-estimates the manager salary by roughly 14.1%.[19]

These misperceptions about peer and manager salaries are similar across different segments of the population, such as by gender, tenure, and payband.[20] Probably because of the sensitive nature of the exercise, we are unaware of other studies that can assess the accuracy of salary perceptions inside a corporation. A notable exception is Lawler (1965), who collected survey data on 326 employees from four privately owned U.S. companies. His findings are qualitatively consistent with ours – for example, he finds that employees tend to systematically under-estimate the salary of their superiors but he does not find such under-estimation for peer salary.

One potential explanation for these misperceptions is that individuals have little information besides their own salary history. We know that most individuals do not report exactly their own salary as their guess for average peer salary: only 35% of them report a guess for average peer salary within 5% of their own salary. However, whatever extra information they use, it does not seem to improve their accuracy: if individuals reported their own salary as guess for the average peer salary, the mean absolute error would be 11.4% (vs. 11.5% in reality), and the bias would be -0.4% (vs. 2.5% in reality).

Extrapolating from own salary to the average manager salary is substantially more chal-

[19] One potential explanation for this under-estimation is that employees may be projecting their own salary forward using their past salary growth linearly instead of exponentially.

[20] The heterogeneity analysis is reported in Appendix A.3.

lenging than extrapolating from own salary to average peer salary – the employee must project her own salary forward, as though being promoted to the managerial position.[21] This circumstance could explain why employees fare worse at guessing manager salaries (mean absolute error of 28%) than peer salaries (mean absolute error of 11.5%).

5.2 Willingness to Pay for Salary Information

To assess whether employees are self-aware of their misperceptions, we use data on the probability distribution of beliefs. This data suggest that, even though they are somewhat optimistic, individuals are mostly aware of how inaccurate their beliefs about peer salaries are. The average individual thinks that there is a 75% probability that their guess for average peer salary falls within 10% of the truth, while the actual share of guesses falling that close to the truth is 55%. In comparison, individuals are largely over-optimistic about the accuracy of their guesses of manager salary. They expect a 74% probability of guessing within 10% of the truth, but the actual share of guesses in this neighborhood is just 24%.

Individuals having misperceptions and being aware of them could have one of the following two explanations. Individuals may not care about these peer and manager salaries, and therefore they choose not to incur costs to acquire information. Alternatively, individuals may care a lot about the salaries, but they don't know where to acquire the information or the acquisition costs are too high. The results from the willingness-to-pay exercise can help to distinguish between these two sources of misperceptions.

When buying information about peer salary, the majority (85%) of respondents made selections that are consistent across scenarios (i.e., their demand functions are non-increasing in price).[22] Following the standard practices (Andersen et al., 2006), the following results focus on subjects with consistent responses – the results are similar under alternative approaches.

Figure 2.a shows the distribution of willingness to pay for the signal about peer salary (i.e., the average salary among a sample of five peers). This willingness to pay has substantial dispersion. On the one hand, some individuals are willing to pay next to nothing for this information: the bottom-25% is willing to pay less than \$6.5, which is less than an hourâs worth of salary.[23] On the other hand, some employees are willing to pay a significant amounts for the information: the top-24% is willing to pay more than \$652, which for most employees

[21] While own salary provides a reasonable guess for average peer salary, it would be a poor guess for average manager salary because manager salaries are systematically higher. Indeed, very few subjects provide a guess of manager salary that is close to their own salary.

[22] For the manager salary, the corresponding figure was 80%. These rates are in line with other studies employing similar methods; for example, in Fuster et al. (2018), 95% of respondents provided consistent responses across scenarios.

[23] To avoid revealing sensitive information about the distribution of pay in the firm, we refrain from providing more precise information.

constitutes over a weeksâ worth of salary. Following the usual methods, we estimate a mean WTP of around $338.[24]

Figure 2.b shows the distribution of willingness to pay for the signal about manager salary. The distribution of WTP seems quite similar for manager salary (Figure 2.b) than for peer salary (Figure 2.a). Using the same methodology, we estimate a mean WTP of about $328. Even though their aggregate distribution is similar, Figure 2.c shows that the WTP values for peer and manager salaries are not perfectly correlated at the individual level (correlation coefficient of 0.28, p-value<0.01). That is, some individuals value the peer information but not the manager information, and vice versa.

To illustrate how large these valuations are, it is useful to compare our results to those from other studies that elicit willingness to pay for other types of information. Relative to the mean WTP for peer and manager information ($328 and $338), these other studies find valuations that are orders of magnitude smaller: the average WTP is $0.40 for travel information (Khattak et al., 2003), $0.80 for food certification information (Angulo et al., 2005), $3 for home energy reports (Allcott and Kessler, 2015), and $4.75 for housing price information (Fuster et al., 2018).[25]

One potential concern is that our estimates of willingness to pay may be sensitive to the elicitation method – for example, the lists of prices given in the hypothetical scenarios may act as a signal for what the employees "should" pay for the information. However, evidence from a follow-up study (Cullen and Perez-Truglia, 2018) suggests that the methodology is not a source for concern: when measuring willingness to pay for information about peer salary using the open-ended method, instead of the multiple price list method, we find a similar distribution of valuations. Indeed, this finding is consistent with prior evidence that measures of willingness to pay are largely similar across different elicitation methods (Brebner and Sonnemans, 2018).

The substantial willingness to pay for salary information suggests that a great deal of the misperceptions arise because acquiring information is difficult rather than because individuals are disinterested. Additionally, the large valuations for some individuals suggest that they may need the information to make high-stakes decisions such as whether to take an outside job offer or request a raise or a promotion. For instance, suppose that you want to buy information about peer salary for use in your salary negotiations. If the information is expected to translate into a salary increase of just 5% and for just one year, then you should be willing to pay more than two weeksâ worth of salary for this information (Stigler, 1962).

[24] This method (Andersen et al., 2006) assumes that the average of the WTP inside each bin is equal to the midpoint of the bin (and for the highest bin, which has no upper bound, the upper bound is set equal to twice the value of the lower bound).

[25] All these amounts were converted to 2017 USD PPP to be comparable with our estimates.

Employees may plan to use the information for a variety of decisions, so the value of the information can add up quite rapidly across all of these different margins.

5.3 Learning

Given that employees had inaccurate beliefs and they were aware of their inaccuracies, we should expect them to learn significantly from the signals provided in the information-provision experiment. In this section, we estimate the simple Bayesian learning model from Section 2.2.

Figure 3 presents a binned scatterplot representation of the Bayesian learning model (equation (4) from section 2.2). Panel (a) corresponds to peer salary, while panel (b) corresponds to manager salary. The y-axis corresponds to the belief updating (i.e., the difference between the posterior and prior belief), while the x-axis corresponds to the information treatment.[26] If individuals incorporated the feedback into their prior beliefs, we should observe a positive association in Figure 3; that is, respondents who overestimated salaries would revise their beliefs downwards when shown the signal, while those who underestimated salaries would revise their beliefs upwards when shown the signal. Moreover, the slope of this regression provides a direct estimate of the learning rate (i.e., the weight that the average subject attaches to the signal provided, relative to the weight attached to the prior belief).[27] Note that the learning rates are expected to be below 1 because the signal provided by the experimenter is based on a sample of five salaries and thus it is subject to sampling variation.

Figures 3.a and 3.b indicate that individuals reacted to the information provision significantly and as predicted by Bayesian learning – indeed, the fit of the linear relationship is almost perfect. According to the slope reported in Figure 3.a, when forming posterior beliefs about peer salary, employees put a weight of 0.51 (SE 0.06) on the signal provided by the experimenter and the remaining weight of 0.49 on their prior belief. In turn, the slope from Figure 3.b suggests that, when forming posterior beliefs about manager salary, employees put a weight of 0.69 (SE 0.03) on the signal provided by the experimenter and the remaining 0.31 on their prior beliefs.[28] The fact that employees gave substantial weight to the information provided by the experimenter is consistent with the prior evidence that, according to the subjective probability beliefs, individuals are not fully confident about their prior beliefs.

[26]As shown in equation (4), this is the difference between the prior belief and the signal interacted by a treatment dummy that indicates if the individual was shown that signal – this regression controls for the difference between the prior belief and the signal without the treatment interaction.

[27]This regression already weeds out spurious reversion to the signal by controlling for the difference between the prior belief and the signal without the treatment interaction. For more details, see Appendix A.4.

[28]Additionally, we find that learning was compartmentalized. Individuals did not use the feedback about peer salary to form beliefs about manager salary – results reported in Appendix A.4.

Additionally, the learning rate being lower for peer salary than for manager salary (0.51 vs. 0.69, p-value<0.001) suggests that individuals had stronger priors about the peer salary than about the manager salary.[29]

5.4 Information Diffusion

Even if the firm did not disclose any information about salaries, employees could form accurate beliefs by sharing salary information. For instance, if all individuals in a peer group shared their own salary with each other, everyone in the group could form an exact guess for the average peer salary. One potential explanation for the large misperceptions seen in the data is that there is not much information diffusion. Indeed, our information-provision experiment provides an opportunity for measuring information diffusion. Intuitively, we can measure whether the information provided to an employee affects the employee's own posterior beliefs and whether it also affects the posterior beliefs of other individuals that are close in the employee network.

The regression results are presented in Table 2. The dependent variables in this regression are the degree of misperceptions in posterior beliefs, measured as the percent absolute error.[30] Columns (1)–(4) correspond to misperceptions about average peer salary, while columns (5)–(8) correspond to misperceptions about average manager salary. The right-hand variables of the regressions correspond to information treatments, such as whether the employee received information herself or whether someone close to the employee received information.[31]

In column (1) of Table 2, the regressor *Received Own* is a dummy for whether the individual was randomly chosen to receive a signal about peer salary. If individuals incorporated the feedback in their posterior beliefs, the treatment assignment should reduce misperceptions. Indeed, receiving the information about peer salary reduces the absolute error by 4.4 percentage points. This effect is not only highly statistically significant (p-value<0.001), but also large in magnitude – it accounts for almost half of the mean of the dependent variable (8.9 percentage points). This finding is consistent with Figure 3, which shows that individuals incorporated the accurate feedback into their posterior beliefs.

For the first test of information diffusion, we try to identify a peer who works in close interaction with the employee and thus may be most likely to share salary information. We define "closest peer" as the peer who has the highest total of emails exchanged (sent

[29]This evidence is also consistent with the view that individuals have little information besides their own salary history and thus they are better equipped to guess the peer salary than to guess the manager salary.

[30]That is, this outcome is the absolute value of the difference between the employee's posterior belief and the truth, divided by the truth.

[31]Additionally, all regressions include the same set of control variables: a linear time trend, the number of peers, and a set of position dummies.

and received) over the three months previous to the start of the experiment. Even though this measure is based on email data, it is plausible that it is also correlated to face to face interactions. To provide supporting evidence, we use the swipe data to calculate a proxy for whether a given pair of employees have lunch together – these are employees who, during lunch time, swipe in and out of the building within 30 seconds of each other. We find that, relative to her other peers, an employee is 53% more likely to grab lunch with the "closest peer".[32]

Column (2) of Table 2 introduces *Closest Peer Received* as additional regressor. This dummy variable indicates whether the respondent's closest peer had already received a signal about peer salary by the time the respondent started the survey.[33] If employees always share the information received in the survey with their closest peer, we would expect this coefficient to be similar to the coefficient on *Received Own*. Instead, we find a coefficient that is close to zero (0.6 percentage points), statistically insignificant, and precisely estimated. We can confidently reject that the coefficient on *Closest Peer Received* is equal to the coefficient on *Received Own* (p-value<0.001). In other words, when we provide information about peer salary to one employee, that information affects her own subsequent perceptions but does not affect the perceptions of her closest peer.

Columns (3) and (4) of Table 2 show two alternative specifications of information diffusion. *Share of Peers Received* measures the share of peers who received information prior to the respondent, while *No. Peers Received* measures the number of peers who received information prior to the respondent.[34] The results are similar: the information spillovers are close to zero and statistically insignificant. Last, column (5)–(8) reproduce the analysis from columns (1)–(4), except with a focus on beliefs about manager salary rather than beliefs about peer salary. Again, we find no evidence of information diffusion.

We can provide a complementary non-experimental test of information diffusion. The theory and evidence indicate that, in presence of social learning, individuals who are most central in a network are the ones who are best informed (Alatas et al., 2016; Banerjee et al., 2013). We can test this hypothesis in our data by comparing the misperceptions in prior beliefs between individuals with above-median and below-median centrality in the email network.[35]

[32]We can calculate this lunch proxy for employees working in the headquarters offices because those are the employees for which there is available swipe data. Our proxy for the probability of grabbing lunch takes the value 18.4% for the closest peer and 12.0% for the other peers.

[33]The average of *Closest Peer Received* is about 0.15. The source of exogenous variation in this regressor arises from the random assignment to information as well as from the random order in which employees were invited to fill out the survey.

[34]The average of *Share of Peers Received* is 0.12, and the average of *No. Peers Received* is 2.77. These variables have the same types of exogenous variation as *Closest Peer Received*.

[35]To measure centrality, we use the directed network of emails sent by employees over the three months prior to the completion of the first survey. We exclude from this sample the emails directed outside of the

Contrary to the hypothesis of social learning, we find that, if anything, misperceptions increase slightly with network centrality. For peer salary, the mean absolute error is 10.7% for individuals below median eigenvalue centrality vs. 12.3% for individuals above median centrality (p-value of difference=0.002); for perceived manager salary, the corresponding values are 28.0% vs. 28.2% (p-value=0.91).

Additionally, we consider a complementary test. If social learning is effective, then individuals who discuss salaries with coworkers should have lower misperceptions. To the contrary, we find misperceptions to be statistically indistinguishable between employees who gossip about salaries and those who do not. For peer salary, the mean absolute errors of prior beliefs are 11.5% for individuals who gossip vs. 11.6% individuals who do not gossip (p-value=0.88); for manager salary, the corresponding values are 27.8% vs. 28.5% (p-value=0.49). This evidence suggests that, even if employees may sometimes discuss salaries with coworkers, they may be sharing noisy or misleading information.

At least two potential explanations may exist for the lack of significant information diffusion. First, employees may possibly see salary information as a rivalrous asset; for example, if one employee shares salary information with a peer, this peer may use the information to get a raise that the first employee could have gotten. Second, employees may prefer to conceal or misrepresent their own salary because they care directly about what others think of them; for example, an employee may feel ashamed to reveal that she gets paid less than a peer. Indeed, in a follow-up study (Cullen and Perez-Truglia, 2018), we present evidence suggestive of these two possibilities.

6 Results: The Effects of Perceived Peer and Manager Salaries on Behavior

In the previous section, we presented evidence about the formation of beliefs about the salaries of others. In this section, we study the effects of those perceptions on the employee's behavior.

6.1 Main Results: Effects on Effort and Performance

Table 3 presents the causal effects of perceived peer and manager salaries on behavior. These effects are estimated with the instrumental variables model outlined in section 2.2. Each column of Table 3 uses a different form of behavior as the dependent variable. The main

institution and emails received from outsiders. These results are based on eigenvalue centrality, but the findings are similar with alternative definitions of centrality.

outcomes of interest are effort and performance, which are presented in columns (1) through (3). Since the right-hand and left-hand variables are defined in logs, the coefficients can be interpreted directly as behavioral elasticities.

Column (1) of Table 3 corresponds to our first measure of effort: the average number of hours worked in the days from the completion of the survey until 90 days later. Recall that this measure is only available for 29% of the sample (i.e., employees based in headquarter offices). The coefficient on *Log(Peer-Salary)* is negative (-0.943) and statistically significant at the 5% level. This coefficient is also economically substantial, implying a peer elasticity of -0.943; that is, increasing the perceived peer salary by 1% would decrease the hours worked by nearly 0.943%. However, the coefficient on *Log(Manager-Salary)* is positive (0.150), statistically significant at the 5% level, and economically significant. This coefficient implies a behavioral peer elasticity of 0.150; that is, increasing the perceived manager salary by 1% would increase the number of hours worked by 0.150%. Note that the manager elasticity is more precisely estimated than peer elasticity because the information-provision experiment induced larger shocks to beliefs about manager salary than to beliefs about peer salary. The main reason for this outcome is that the prior misperceptions about manager salary are more substantial than the prior misperceptions about peer salary; therefore, there is more room to shift beliefs by correcting misperceptions.[36]

Column (2) of Table 3 uses our alternative measure of effort: the average number of emails sent, which is available for the entire subject pool. The peer elasticity is negative (-0.431) and significant at the 5% level, and the manager elasticity is positive (0.130) and significant at the 1% level. The elasticities from column (1) are also quantitatively consistent with the elasticities from column (2). We cannot reject that the peer elasticity for hours worked (-0.943) is equal to the peer elasticity for emails (-0.431) – p-value=0.271; and we cannot reject that the manager elasticity for hours worked (0.150) is equal to the manager elasticity for emails sent (0.130) – p-value=0.816.

Column (3) of Table 3 uses as dependent variable our only measure of performance: the sales performance index. This outcome is available only for 38.4% of employees, that is, those who have a sales role. Again, the peer elasticity is negative (-0.731), statistically significant at the 5% level, and on the same order of magnitude as the effort elasticities. The manager elasticity is positive (0.106) and on the same order of magnitude as the effort elasticities, but it is slightly smaller and less precisely estimated; thus, it becomes statistically insignificant (p-value=0.383).

One of the most important and robust findings is that the peer and manager elastici-

[36] Additionally, given a level of misperceptions, individuals reacted more to feedback about manager salary than to feedback about peer salary – presumably because they had a weaker prior belief about manager salaries.

ties have opposite signs. A higher perceived peer salary has negative effects on effort and performance, while a higher perceived manager salary has positive effects on effort and performance. To provide a more rigorous comparison, the bottom of each column of Table 3 reports the p-value of the test of the null hypothesis that the peer-elasticity is equal to the manager-elasticity. We always reject this null hypothesis, with p-values of 0.026 for hours worked, 0.007 for emails sent, and <0.001 for sales performance.

6.2 Effects on Career Outcomes

One unique aspect of our setting is that subjects are in a continuing contract with the firm, and thus we can follow what happens to this relationship going forward, such as through exits or salary negotiations.

The effects on these career outcomes are reported in columns (4)–(7) of Table 3. Columns (4) and (5) explore two forms of retention. Column (4) uses a dummy variable for whether the employee leaves the firm as the dependent variable. The results suggest that a 1% increase in perceived peer salary increases the probability of leaving the company by 0.235 percentage points, which is statistically significant at the 5% level. Even though this elasticity is not as large as the corresponding elasticities for effort and performance, its direction is consistent: a higher perceived peer salary demotivates employees to the extent that they are more likely to leave the firm. With regard to vertical comparisons, a 1% increase in perceived manager salary decreases the probability of leaving the company by 0.015 percentage points, but the effect is economically and statistically insignificant. In column (5), the dependent variable is a dummy indicating whether the individual is transferred to another unit within the same firm. Even though the signs of the cross-salary elasticities are consistent with those from column (4), the coefficients are closer to zero and statistically insignificant. In column (6), the dependent variable is the logarithm of the base salary three months after the completion of the survey. Both the peer and manager elasticities are close to zero, statistically insignificant, and precisely estimated. These results imply that changes in salary perceptions did not materialize into salary negotiations in just 3 months. Similarly, column (8) uses a dummy for change in position title as dependent variable. Again, both the peer and manager elasticities are close to zero and statistically insignificant.

In summary, the evidence suggests that perceived peer and manager salaries have a large effect on effort and performance, but they do not affect career outcomes. The only exception is a significant peer elasticity on the retention margin. However, it is important to note that these perceptions could possibly affect career outcomes with longer horizons, such as years into the future – we return to this discussion in section 6.4, where we present the effects on salary expectations.

6.3 Robustness Checks

The third and fourth rows of Table 3 present the event-study falsification tests. These coefficients are estimated using the same specification as the first two rows, but using pre-treatment outcomes instead of post-treatment outcomes as the dependent variables. Intuitively, we expect these coefficients to be close to zero and statistically insignificant, because the information that was randomly provided on the date of the survey could not have possibly affected the behavior prior to the survey date. The estimates are consistent with this robustness check. For instance, the pre-treatment coefficients from column (1) are close to zero (-0.205 for peer elasticity and 0.001 for manager elasticity) and statistically insignificant at conventional levels.

When using instrumental variables regressions, one potential concern pertains to weak instruments. According to the evidence from Section 5.3, the information provision experiment has large effects on the posterior beliefs, which is suggestive of a strong first stage. For a more formal assessment of weak instruments, the bottom of Table 3 reports the Cragg-Donald F statistic. In column (1), this statistic takes a value of 29.8, which is substantially above the rule of thumb used to detect weak instruments (Stock and Yogo, 2005). Moreover, this statistics is even larger in the rest of the specifications: 98.2 in column (3) and over 200 in each of the five remaining columns.

We also test the sensitivity of the results to the event window. Table 4 presents these results. The first pair of rows correspond to the baseline specification, looking at the effects 90 days after the survey completion – by construction, they are identical to the first pair of rows from Table 3. In the second pair of rows, we use an alternative window: 180 days after the survey completion. The results indicate that the effects are less precisely estimated in the longer time window, which is to be expected, given that the outcomes become less predictable as we move further into the future. Regarding peer elasticity, the results are qualitatively and quantitatively similar across the two specifications. Regarding the manager elasticity, the coefficients are quantitatively similar, although there are some differences in statistical significance.[37] The bottom rows of Table 4 report difference tests between the elasticities computed in the 90-day and 180-day windows. The results are statistically indistinguishable across the two specifications, and we cannot reject the null hypothesis of equal coefficients in any of the 14 tests. However, because of the precision of the estimates, we cannot rule out that the effects diminished somewhat over time.

In Appendix A.5, we explore a number of additional tests. First, we split the cross-salary

[37] Two of the coefficients (for hours worked and number of emails) that are statistically significant in the baseline specification become borderline insignificant in the longer time period, but one of the coefficients (for sales performance) that was statistically insignificant in the baseline specification becomes highly significant with the longer time period.

elasticities by subgroups of the population, such as females versus males or higher versus lower pay bands. We do not find any statistically significant evidence of heterogeneity. We also relax some of the assumptions from the baseline model. First, the model assumes a (log-log) linear relationship between perceived salaries and behavior. Using binned scatterplots, we show that this linear assumption is a reasonable approximation. Second, our baseline model assumes symmetric effects. If there were assymetries, that could lead to an under-estimation of the importance of salary comparisons. We find that, for the exit outcome, the peer elasticity is asymmetrical with respect to whether the own salary is below or above the peer average. This result is consistent with the asymmetry reported in Card et al. (2012) for a similar outcome (stated intent to find a new job). However, we do not find any robust evidence of asymmetric effects for the main outcomes (effort and performance).

We must note that the regressions for instrumental variables identify local average treatment effects – that is, weighted averages of the elasticities, with a higher weight given to employees whose beliefs are more affected by the information-provision experiment. By construction, this weight will be higher for individuals who have larger prior misperceptions and, conditional on the misperceptions, for individuals who react more to feedback. We provide evidence that the degree of misperceptions in prior beliefs and the learning rates are largely unrelated to observable characteristics such as gender, tenure, and occupation.[38] As a result, our estimated elasticities are plausibly quite representative of the entire subject pool.

6.4 Preferred Interpretation of the Evidence

In this section, we discuss some potential interpretations of the findings and provide suggestive evidence for them. While several potential mechanisms may be used to rationalize the cross-salary elasticities, we focus on two mechanisms that have played a prominent role in the theoretical and empirical literatures: expected future salary and social comparisons.[39] Indeed, the interplay between these two mechanisms plays a central role in some models of employee compensation (Lazear, 1989; Ederer and Patacconi, 2010).

The first mechanism, salary expectations, posits that employees use the salaries of peers and managers to form expectations about their own future salary (Lazear and Rosen, 1981; Gibbons and Murphy, 1992; Holmstrom, 1999). An individual who discovers that peers are getting paid more may see a positive signal of own future salary, as in the "tunnel effect" (Hirschman and Rothschild, 1973). This change in expectations should make working hard

[38] Results reported in Appendix A.4 and A.7, respectively.

[39] For example, these two forces are used to explain the effects of relative income on subjective well-being (Senik, 2004) and job satisfaction (Godechot and Senik, 2015; Clark et al., 2008). These papers argue that the effect of relative income depends on the group of reference – for some groups the negative "comparison" effect dominates, while for other groups the positive "information" effect dominates.

to keep the current job at the firm more attractive. Similarly, an increase in the perceived manager salary could also increase the expected own future salary, if the employee expects to be promoted to the managerial position. This expectation should also incentivize the individual to work harder in order to get promoted. As a result, this salary expectations mechanism predicts $\eta_{peer} > 0$ and $\eta_{mgr} > 0$. The second mechanism, social comparisons, posits that employee morale depends on employees' compensation relative to other employees (Frank, 1984; Romer, 1984; Summers, 1988; Akerlof and Yellen, 1990; Lazear, 1989) – indeed, a similar argument is used to explain wage rigidities (Solow, 1979; Bewley, 1999). Holding own salary constant, a higher perceived peer salary (or a higher perceived manager salary) worsens the employee's relative standing and thus leads to lower effort.[40] As a result, this mechanism predicts $\eta_{peer} < 0$ and $\eta_{mgr} < 0$.

The signs of the elasticities contain some suggestive information about which of these two mechanisms dominates. The negative sign of the peer-elasticity suggests that the social comparisons channel dominates over the salary expectations channel. In contrast, the positive sign of the manager-elasticity suggests that the salary expectations channel dominates over the social comparisons channel.

To provide suggestive evidence about these and other mechanisms, Table 5 reports the effects of salary perceptions on the different survey outcomes. Each of the six columns corresponds to a different dependent variable, based on the six questions included in the end of the survey. Some of these dependent variables are not expressed in logs or in percentage points, and they should be interpreted as semi-elasticities instead of elasticities. All coefficients are estimated with the same instrumental variables specification used for the behavioral outcomes, except that we do not observe pre-treatment survey outcomes, so we cannot use them for falsification tests or include them as control variables. We include the following additional control variables as substitutes: dummies for sales role, pay band, unit, and position title.

Column (1) and (2) of Table 5 uses pay satisfaction and job satisfaction as dependent variables. Finding effects on these outcomes would suggest that the behavioral effects operate through employee morale (Clark and Oswald, 1996). In column (1), the coefficient on peer salary (-0.762) is negative and statistically significant at the 10% level. This effect is economically large, implying that a 1% higher peer salary decreases pay satisfaction by roughly 0.83% of a standard deviation. In contrast, the coefficient on perceived manager salary is close to zero (-0.015), statistically insignificant, and precisely estimated. Indeed, we can reject the null hypothesis that the coefficients for peer salary and manager salary are equal (p-value=0.084). The results for job satisfaction are somewhat consistent with the

[40] These preferences can be thought as being based on the combination of relative deprivation (Frank, 1984; Luttmer, 2005) and reciprocity (Akerlof, 1982).

results for pay satisfaction. The coefficient on peer salary changes from -0.762 in column (1) to -0.444 in column (2), which is consistent with the expectation that salary is a more important determinant of pay satisfaction than of overall job satisfaction. Even though the peer coefficient becomes statistically insignificant in column (2), we cannot rule out large negative effects. Consistent with the results from column (1), the coefficient on peer salary is close to zero (-0.086), statistically insignificant, and precisely estimated. Column (3) provides a complementary test of the social comparisons channel by using preferences for within-firm redistribution as dependent variable. The coefficient on peer salary (0.373) is statistically significant at the 10% level and economically significant: increasing the perceived peer salary by 1% increases preferences for redistribution by 0.65% of a standard deviation. In comparison, the coefficient on manager salary is close to zero (0.008), statistically insignificant, and precisely estimated. In summary, the evidence from columns (1) through (3) suggests that social comparisons are a mediating factor for the peer elasticities but play no role in the manager elasticities.

Next, we provide a test of the salary expectations channel. Columns (4) and (5) of Table 5 measure the effects of salary expectations on the (log) expected future salaries 1 year and 5 years ahead, respectively. First, we focus on the effects of vertical comparisons. Column (4) suggests an elasticity between perceived manager salary and 1-year-ahead expected own salary that is close to zero (0.025), statistically insignificant, and precisely estimated. This null effect should not be surprising given that employees would almost never be promoted to their manager's position in just one year. In turn, column (5) shows that the effect of perceived manager salary on 5-year-ahead salary is positive (0.166), precisely estimated, and statistically significant at the 1% level. Indeed, we can demonstrate that, under a reasonable assumption, the magnitude of this elasticity (0.166) can fully account for the magnitude of the elasticity between hours worked and manager salary (0.150). Due to the 0.166 elasticity, a 1% increase in perceived manager salary would increase expected own salary by 0.166%. Now assume an elasticity between hours worked and future salary of 0.9. That would imply that the 0.166% increase in expected own salary translates into a 0.150% ($= 0.166 \cdot 0.9$) increase in hours worked.

Second, we focus on the effects of horizontal comparisons on expected future salary. The peer elasticity is 0.071 with respect to 1-year-ahead salary (column (4) of Table 5) and 0.280 with respect to 5-year-ahead salary (column (5)). Although these two coefficients are statistically insignificant at conventional levels (p-values of 0.431 and 0.111, respectively), they are on the same order of magnitude as the corresponding coefficients for perceived manager salary. This outcome is suggestive, although weaker, evidence that employees become more optimistic about their own future salary when they find out that their peers earn more – just like in the "tunnel effect" (Hirschman and Rothschild, 1973).

In summary, the findings from columns (4) and (5) of Table 5 suggest that the salary expectations channel is present both in horizontal and vertical comparisons. Additionally, these findings provide suggestive evidence for the potential explanations of why the employees are willing to pay substantial amounts for the salary information: this information may allow them to pursue career paths with higher future salaries, and they may plan to use the salary information in their future salary negotiations.

To provide further evidence on the mechanisms underlying the manager elasticities, we can exploit variation in the distance between the employee's position and the managerial position.[41] Consider the case of an employee learning about the salary of a manager that is several promotions away – that is, a position she will probably not be promoted to. In that case, we would not expect the employee to extrapolate the information about the manager salary to her own salary expectations, or at least not to a great extent. As a result, the corresponding manager elasticity should be either zero – if the social comparisons are absent – or negative – if the social comparisons are at play.

The results from the heterogeneity analysis are presented in Table 6. The manager elasticities are broken down into two groups: *Closer* corresponds to managerial positions that are within the reach of the individual, while *Farther* corresponds to managerial positions that are outside the reach of the individual. We present results from two models, based on the two survey measures of the distance to the manager. In the first panel (Model 1), *Closer* and *Farther* are defined based on the subjective probability of reaching the managerial position. In the second panel (Model 2), these groups are defined based on the expected numbers of promotions needed to reach the managerial position.

In columns (1) through (3), the dependent variables are the three measures of effort and performance. The results suggest that for hours worked and emails sent (columns (1) and (2)), and regardless of the measure of distance that we use (Models 1 and 2), the manager elasticity is positive and significant when the employee is closer to the managerial position, but close to zero and insignificant when the employee is farther from the managerial position. For instance, in Model 1 and column (1), the manager elasticity is 0.212 for the closer managers and -0.074 for the ones farther away, with a p-value of the difference of 0.040. However, the rest of the differences, while economically substantial, are imprecisely estimated and thus statistically insignificant (p-values of 0.170, 0.243, and 0.212). The elasticities for sales performance (column (3)) are quite imprecisely estimated for the managers that are farther away, and thus the results are uninformative. In columns (4) and (5), the dependent variables are the survey measures of expected future salary. Column (4) indicates that

[41] This distance was not randomized, and thus the results should be interpreted with caution.

the perceived manager salary does not affect expectations about the 1-year-ahead salary, regardless of the distance to the manager. Column (5) suggests that the perceived manager salary affects the 5-year-ahead salary expectations, and it does so more strongly when the managerial position is closer than when it is farther away. However, this result must be taken with a grain of salt because the differences are statistically insignificant (p-values of 0.229 and 0.560). In summary, the evidence from Table 6 makes two points. First, because of the higher effects for managerial positions that are closer, it provides supportive evidence that the salary expectations channel is present for vertical comparisons. Second, because of the lack of negative effects in positions that are farther away, it provides further suggestive evidence that the social comparisons channel is not at play in the vertical comparisons.

The evidence suggests that social comparisons are present for horizontal comparisons but not for vertical comparisons. At least two potential explanations exist for this difference. First, as has been suggested in the broader literature on concerns for relative standing (Clark and Senik, 2010), employees may only care about their standing in an specific reference group – in this case, their peers. The second possibility is that employees feel demoralized about horizontal comparisons because, given that everyone has the same powers and responsibilities, they perceive these salary differences as unfair. Instead, employees may find it easier to justify vertical inequality – for instance, they may think that the manager deserves the higher salary because she adds more value to the firm or because she has to deal with more stress. Indeed, this finding is also related to a literature stream about preferences for redistribution, according to which some poor people do not want to tax the rich if they think the rich are deserving of their wealth (Di Tella et al., 2016).

6.5 Alternative Interpretations of the Evidence

In this section, we discuss some alternative mechanisms that could potentially explain our findings.

One possible explanation for the peer elasticities is that, in the spirit of career concerns, employees use their relative salary to infer their relative productivity.[42] Column (6) of Table 5 provides a test of this hypothesis, by looking at the effects of peer and manager salaries on perceived productivity rank. This outcome, which is based on an incentivized question, can take values from 0 (least productive) to 1 (most productive). The coefficients on peer and manager salaries are close to zero, precisely estimated, and statistically insignificant: a 1% increase in peer salary decreases perceived productivity rank by just 0.044 percentage points, and a 1% increase in manager salary increases perceived productivity rank by less than 0.001

[42]Relatedly, experimental evidence suggests that employees can react to the effort or productivity of other employees (Mas and Moretti, 2009; Bandiera et al., 2010).

percentage points. This constitutes evidence that the effect of salary information does not operate through changes in beliefs about relative productivity.[43]

A related explanation for the effects of perceived peer salary is that individuals use those perceptions to form beliefs about the returns to effort. If the employee discovers that her job is less meritocratic than previously thought, she may feel disinclined to exert effort. However, that mechanism does not fit the results because it predicts a peer elasticity of zero. Intuitively, when getting a signal that their peers are doing better, the least productive individuals should infer that the returns on effort are higher, while the most productive individuals should infer that the returns on effort are lower. Moreover, using the survey data on perceived productivity rank, we do not find any evidence consistent with this type of heterogeneity.[44]

It is possible that employees dislike disparities in salaries regardless of where they stand in the salary distribution (Breza et al., 2018). Our information interventions were designed to shift beliefs about the relative standing rather than beliefs on the dispersion of salaries. If employees process the information rationally, then a signal about the average peer salary should have roughly no effect on the belief about the dispersion of salaries within the peer group (Hoff, 2009). As a result, it is unlikely that the effects of average peer salary operate through perceptions of peer inequity. In other words, while inequity aversion may be significant, we would need a different experiment to measure it.

Last, the peer elasticity could be the product of employees using the peer salary information to form beliefs about the salary that they could earn working for another firm (Shapiro and Stiglitz, 1984). While this mechanism can provide a straightforward explanation for the effects on employee retention, however, it does not provide a straightforward explanation for the negative effects on effort and performance.[45]

7 Conclusions

We presented evidence from a field experiment involving 2,060 employees from a multi-billion dollar corporation. The research design combines survey data, administrative data, and an

[43]To the extent that these productivity ratings are assigned by managers, this evidence also goes against the related mechanism that employees use their peer relative salary to infer what their managers think of them. There are other pieces of suggestive evidence against this channel. First, employees probably do not need to make this type of inferences because they get plenty of feedback from their managers, such as through their annual reviews. Second, many factors influences the employee's salary and are outside the manager's control.

[44]Results presented in Appendix A.7.

[45]For example, employees may want to work harder if they want the current firm to respond to outside offers. Similarly, employees may want to work harder if they need the current manager to recommend them for the outside jobs.

information-provision experiment to shed light on how employees learn about the salaries of their peers and managers and how their beliefs about those salaries affect their behavior.

We documented large misperceptions about the salaries of peers and managers and identified some of the sources of these misperceptions. Additionally, we showed that perceptions about the salaries of peers and managers have significant effects on employee effort, performance, and retention. We suggested mechanisms underlying these cross-salary elasticities. In horizontal comparisons, both the social comparisons and salary expectations channel are present, but the social concerns dominate; in vertical comparisons, we find suggestive evidence of the salary expectations channel, but no such evidence of social comparisons.

Our findings have some implications for understanding how firms operate. We find that rewarding one employee with a higher salary has a negative externality on the effort of all peers. In contrast, increasing the salary of a group of managers has a positive externality on the behavior of all subordinates. Because of these externalities, firms may find it optimal to load rewards vertically rather than horizontally. That is, firms may want to motivate employees with the prospect of a higher salary upon promotion rather than through performance pay. As a result, our evidence can explain why, as documented by Baker et al. (1988), firms provide most of the financial incentives vertically instead of horizontally.[46]

Similarly, this evidence may shed light on why employees tolerate pay discrimination, such as the gender-based wage gap. Our evidence suggesting that vertical discrimination may be less discouraging than horizontal discrimination could explain why the bulk of the gender pay gap is loaded vertically rather than horizontally. For instance, in the firm where the experiment was conducted, 92% of the gender pay gap comes from vertical differences and only 8% through horizontal differences – a similar decomposition has been found in firms in other countries such as the United States (Barth et al., 2017).

Last, the view that social comparisons put pressure to compress salary differentials within the firm is widespread (Frank, 1984). Our evidence suggests that this view is only true in a narrow sense. While this channel may force firms to reduce horizontal inequality, they are not restricted in their use of vertical inequality. Moreover, academics and policy makers have proposed pay transparency policies with the intent of reducing pay inequality. To the extent that employees are not bothered by vertical comparisons, our evidence suggests that these policies may not be as effective as previously thought.

[46] For a rigorous analysis of the co-evolution of compensation schemes and social concerns, see MacLeod (2007) and MacLeod and Malcomson (1998).

References

Akerlof, G. A. (1982). Labor Contracts as Partial Gift Exchang. *Quarterly Journal of Economics* (97), 543–569.

Akerlof, G. A. and J. L. Yellen (1990). The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics 105*(2), 255–283.

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016). Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia. *American Economic Review 106*(7), 1663–1704.

Allcott, H. and J. B. Kessler (2015). The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons. Working Paper 21671, National Bureau of Economic Research.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006). Elicitation using multiple price list formats. *Experimental Economics 9*(4), 383–405.

Angulo, A. M., J. M. Gil, and L. Tamburo (2005). Food Safety and Consumers' Willingness to Pay for Labelled Beef in Spain. *Journal of Food Products Marketing 11*(3), 89–105.

Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar (2016). The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment. *Review of Economics and Statistics 98*(3), 503–523.

Baker, G., M. Jensen, and K. Murphy (1988). Compensation and Incentives: Practice vs. Theory. *The Journal of Finance 43*(3), 593–616.

Bandiera, O., I. Barankay, and I. Rasul (2010). Social Incentives in the Workplace. *The Review of Economic Studies 77*(2), 417–458.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The diffusion of microfinance. *Science 341*(6144), 1236498.

Barth, E., S. Kerr, and C. Olivetti (2017). The dynamics of gender earnings differentials: Evidence from establishment data. *NBER Working Paper No. 23381*.

Bewley, T. (1999). *Why Wages Don't Fall During a Recession*. Harvard University Press.

Bloom, N., F. Guvenen, D. J. Price, and J. Song (2015). Firming Up Inequality. *NBER Working Paper No. 21199*.

Bottan, N. and R. Perez-Truglia (2017). Choosing Your Pond: Revealed-Preference Estimates of Relative Income Concerns. *NBER Working Paper No. 23615.*.

Bracha, A., U. Gneezy, and G. Loewenstein (2015). Relative pay and labor supply. *Journal of Labor Economics 33*(2), 297–315.

Brebner, S. and J. Sonnemans (2018). Does the elicitation method impact the WTA/WTP disparity? *Journal of Behavioral and Experimental Economics 73*(C), 40–45.

Breza, E., S. Kaur, and Y. Shamdasani (2018). The Morale Effects of Pay Inequality. *The Quarterly Journal of Economics*.

Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review 102*(6), 2981–3003.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Charness, G. and P. Kuhn (2007). Does Pay Inequality Affect Worker Effort? Experimental Evidence. *Journal of Labor Economics 25*(4), 693–723.

Clark, A. E., P. Frijters, and M. A. Shields (2008). Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles. *Journal of Economic Literature 46*(1), 95–144.

Clark, A. E., D. Masclet, and M. C. Villeval (2010). Effort and Comparison Income: Experimental and Survey Evidence. *ILR Review 63*(3), 407–426.

Clark, A. E. and A. J. Oswald (1996). Satisfaction and comparison income. *Journal of Public Economics 61*(3), 359–381.

Clark, A. E. and C. Senik (2010). Who Compares to Whom? The Anatomy of Income Comparisons in Europe*. *The Economic Journal 120*(544), 573–594.

Cohn, A., E. Fehr, B. Herrmann, and F. Schneider (2014). Social Comparison and Effort Provision: Evidence from a Field Experiment. *Journal of the European Economic Association 12*.

Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics 98*, 100–112.

Cullen, Z. and B. Pakzad-Hurson (2016). Equilibrium Effects of Pay Transparency in a Simple Labor Market. *Working Paper*.

Cullen, Z. and R. Perez-Truglia (2018). Social Learning with Sensitive Data: Lessons from a Field Experiment.

Di Tella, R., J. Dubra, and A. L. Lagomarsino (2016). Meet the Oligarchs: Business Legitimacy, State Capacity and Taxation. Working Paper 22934, National Bureau of Economic Research.

Dube, A., L. Giuliano, and J. Leonard (2015). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. IZA Discussion Papers 9149.

Ederer, F. and A. Patacconi (2010). Interpersonal comparison, status and ambition in organizations. *Journal of Economic Behavior & Organization 75*(2), 348–363.

Eil, D. and J. M. Rao (2011). The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. *American Economic Journal: Microeconomics 3*(2), 114–138.

Frank, R. (1984). Are Workers Paid Their Marginal Products? *American Economic Review 74*(4), 549–571.

Fuster, A., R. Perez-Truglia, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Working Paper*.

Gächter, S. and C. Thöni (2010). Social comparison and performance: Experimental evidence on the fair wageâeffort hypothesis. *Journal of Economic Behavior & Organization 76*(3), 531–543.

Gibbons, R. and K. J. Murphy (1992). Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence. *Journal of Political Economy 100*(3), 468–505.

Godechot, O. and C. Senik (2015). Wage comparisons in and out of the firm. Evidence from a matched employer-employee French database. *Journal of Economic Behavior and Organization 117*, 395–410.

Hegewisch, A., C. Williams, and R. Drago (2011). Pay Secrecy and Wage Discrimination. *Institute for Women's Policy Research*.

Hirschman, A. O. and M. Rothschild (1973). The Changing Tolerance for Income Inequality in the Course of Economic Development. *The Quarterly Journal of Economics 87*(4), 544–566.

Hoff, P. D. (2009). *A first course in Bayesian statistical methods*. Springer Science & Business Media.

Holmstrom, B. (1999). Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies 66*(1), 169–182.

Huet-Vaughn, E. (2017). Do social comparisons motivate workers? A field experiment on relative earnings, labor supply and the inhibitory effect of pay inequality. *Working Paper*.

Karadja, M., J. Mollerstrom, and D. Seim (2017). Richer (and Holier) Than Thou? The Effect of Relative Income Improvements on Demand for Redistribution. *The Review of Economics and Statistics 99*(2), 201–212.

Khattak, A. J., Y. Yim, and L. S. Prokopy (2003). Willingness to pay for travel information. *Transportation Research Part C: Emerging Technologies 11*(2), 137–159.

Lawler, E. E. (1965). Managers' Perceptions of their Subordinates' Pay and of Their Superiors' Pay. *Personnel Psychology 18*(4), 413–422.

Lazear, E. P. (1989). Pay Equality and Industrial Politics. *Journal of Political Economy 97*(3), pp. 561–580.

Lazear, E. P. and S. Rosen (1981). Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy 89*(5), 841–864.

Luttmer, E. F. P. (2005). Neighbors as Negatives: Relative Earnings and Well-Being. *The Quarterly Journal of Economics 120*(3), 963–1002.

MacLeod, W. B. (2007). Can Contract Theory Explain Social Preferences? *The American Economic Review 97*(2), 187–192.

MacLeod, W. B. and J. M. Malcomson (1998). Motivation and Markets. *The American Economic Review 88*(3), 388–411.

Mas, A. (2017). Does Transparency Lead to Pay Compression? *Journal of Political Economy 125*(5), 1683–1721.

Mas, A. and E. Moretti (2009). Peers at Work. *The American Economic Review 99*(1), pp. 112–145.

McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics 99*(2), 210–221.

Medoff, J. L. and K. G. Abraham (1980). Experience, Performance, and Earnings*. *The Quarterly Journal of Economics 95*(4), 703–736.

Perez-Truglia, R. (2015). The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment. *Working Paper*.

Perez-Truglia, R. and G. Cruces (2017). Partisan interactions: Evidence from a field experiment in the United States. *Journal of Political Economy 125*(4).

Perez-Truglia, R. and U. Troiano (2015). Shaming Tax Delinquents. Working Paper 21264, National Bureau of Economic Research.

Pritchard, R. D., Marvin D. Dunnette, and D. O. Gorgenson (1972). Effects of perception of equity and inequity on worker performance and satisfaction. *Journal of Applied Psychology 56*(1), 75.

Rege, M. and I. Solli (2015). Lagging behind the joneses: The impact of relative earnings on job separation. *Working Paper*.

Romer, D. (1984). The Theory of Social Custom: A Modification and Some Extensions. *The Quarterly Journal of Economics 99*(4), 717.

Schmitt, D. R. and G. Marwell (1972). Withdrawal and reward reallocation as responses to inequity. *Journal of Experimental Social Psychology 8*(3), 207–221.

Senik, C. (2004). When information dominates comparison: Learning from Russian subjective panel data. *Journal of Public Economics 88*(9), 2099–2123.

Shapiro, C. and J. E. Stiglitz (1984). Equilibrium Unemployment as a Worker Discipline Device. *The American Economic Review 74*(3), 433–444.

Solow, R. M. (1979). Another possible source of wage stickiness. *Journal of Macroeconomics 1*(1), 79–82.

Stigler, G. J. (1962). Information in the Labor Market. *Journal of Political Economy 70*(5, Part 2), 94–105.

Stock, J. H. and M. Yogo (2005). *Testing for Weak Instruments in Linear IV Regression*, pp. 80–108. Cambridge University Press.

Summers, L. H. (1988). Relative Wages, Efficiency Wages, and Keynesian Unemployment. *The American Economic Review 78*, 383–388.

Valenzi, E. R. and I. R. Andrews (1971). Effect of hourly overpay and underpay inequity when tested with a new induction procedure. *Journal of Applied Psychology 55*(1), 22–27.

Figure 1: Salary Misperceptions

a. Own Salary

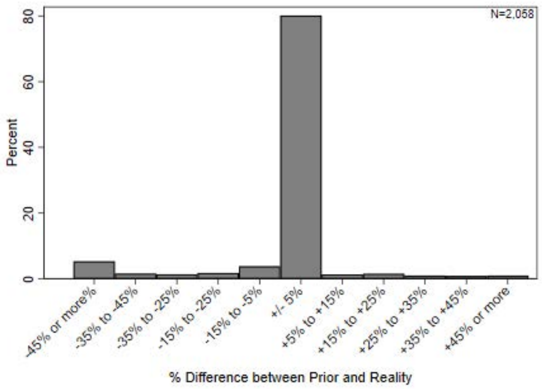

b. Peer Salary

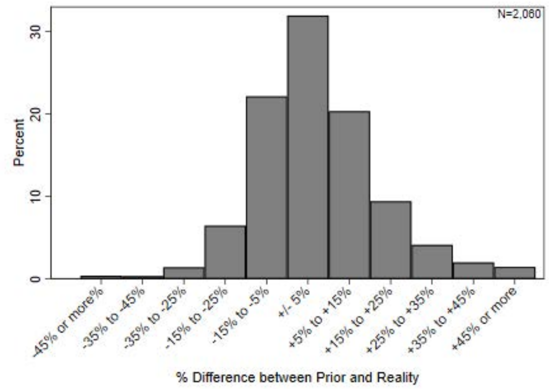

c. Manager Salary

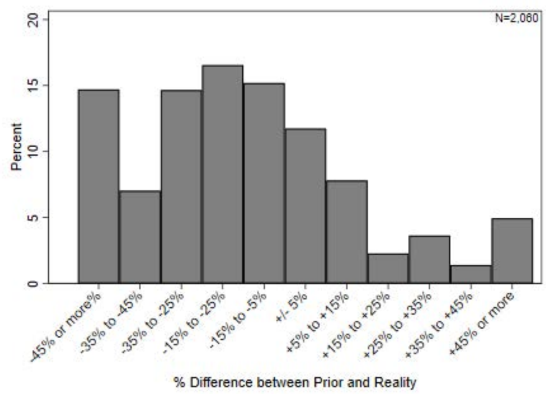

Notes: Histograms of the salary misperceptions, defined as the difference between the employee's prior belief about a particular salary (according an incentivized survey question) and the actual salary (according to the firm's administrative records), divided by the actual salary. Panel (a) corresponds to misperceptions about own salary. Panel (b) corresponds to misperceptions about the average peer salary. Panel (c) corresponds to misperceptions about the average manager salary.

Figure 2: Willingness to Pay for Information about Peer and Manager Salaries

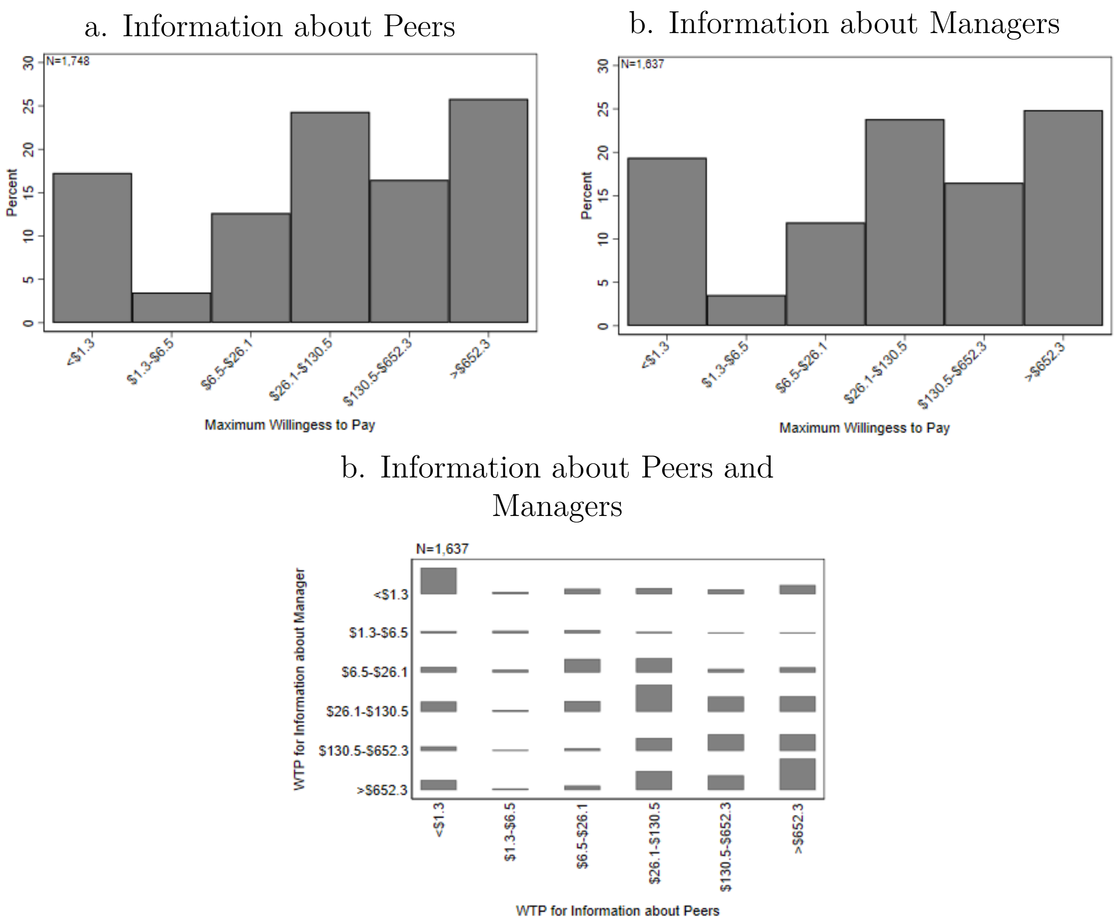

Notes: Histograms of the willingness to pay (WTP) for a specific information piece, based on the responses to multiple price list questions. Panel (a) corresponds to the WTP for the average salary among a random sample of 5 peers. Panel (b) corresponds to the WTP for the average salary among a random sample of 5 managers. Panel (c) corresponds to the joint distribution of (a) and (b). The samples are restricted to the subset of respondents with consistent responses across the five price scenarios.

Figure 3: Learning from Information-Provision Experiment

a. Peer Salary

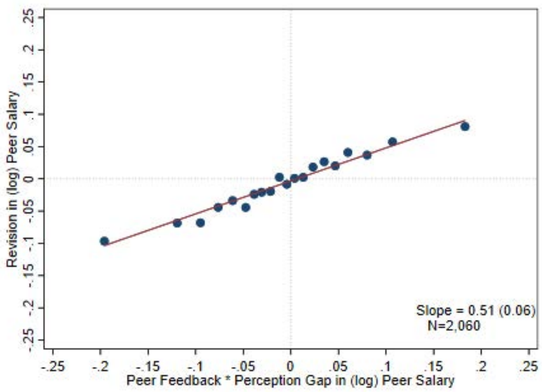

b. Manager Salary

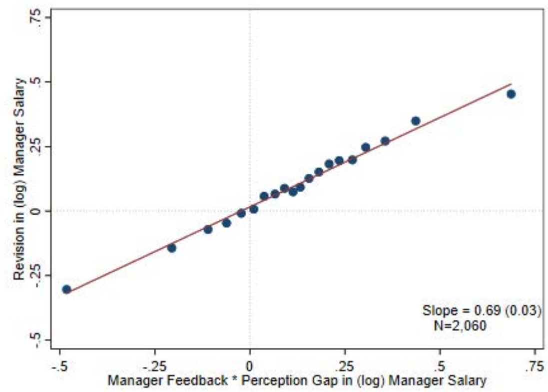

Notes: Partial regression binned scatterplot of the Bayesian learning equation (4) presented in section 2.2. Panel (a) corresponds to learning about the average peer salary, while panel (b) corresponds to learning about the average manager salary. The y-axis corresponds to the respondent's update: i.e., the posterior belief minus the prior belief. The x-axis corresponds to the information treatment: the difference between the feedback chosen for the employee (e.g., the average salary among the random sample of 5 peers) and the employee's prior belief, multiplied by a dummy variable for whether the information was randomly chosen to be shown to the respondent. The regression controls for the difference between the feedback chosen for the employee and the employee's prior belief; also, it controls for the prior belief and position title dummies. The slope comes from a linear regression, with standard errors reported in parentheses (clustered at position level).

Table 1: Descriptive Statistics and Randomization Balance Test

	All	Treatment Group (by Information)				
		Peer & Manager	Peer	Manager	None	P-value
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.73	0.74	0.71	0.74	0.74	0.32
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	
Age	29.20	28.92	29.35	29.35	29.19	0.99
	(0.11)	(0.19)	(0.21)	(0.22)	(0.22)	
College (or Higher)	0.86	0.86	0.87	0.84	0.86	0.14
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	
Tenure (Years)	4.99	4.92	5.08	5.14	4.79	0.81
	(0.08)	(0.14)	(0.15)	(0.16)	(0.16)	
Own Salary (Masked)	0.72	0.72	0.72	0.72	0.72	0.93
	(0.01)	(0.03)	(0.02)	(0.02)	(0.02)	
Avg. Peer Salary (Masked)	0.72	0.73	0.72	0.72	0.73	0.91
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	
Avg. Manager Salary (Masked)	2.84	2.86	2.89	2.80	2.80	0.54
	(0.05)	(0.10)	(0.10)	(0.10)	(0.11)	
Observations	2,060	559	528	510	463	

Notes: Average pre-treatment characteristics of the employees, with standard errors in parentheses. *Female* takes the value 1 if the employee is female and 0 otherwise. *Age* is the employee's age (in years) as of March 2017. *College* takes the value 1 if the employee finished College or a higher degree, and 0 otherwise. *Tenure* is the number of years from the date when the employee joined the company until March 2017. *Own Salary* is the employee base monthly salary as of March 2017. *Avg. Peer Salary* and *Avg. Manager Salary* are the true average and peer salaries. Due to the sensitive nature of the data, we do not reveal the unit of measurement for salary variables. Column (1) corresponds to the entire subject pool, while columns (2) through (5) correspond to the four treatment groups that subjects were randomly assigned to: receiving information about both their peers' and manager's salary (column (2)); receiving information about the average salary of peers but no salary information about their manager (column (3)); receiving information about the salary of their manager but not their peers (column (4)); and receiving no salary information (column (5)).

Table 2: Information Diffusion

	Misperceptions about Peer Salary				Misperceptions about Manager Salary			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Information Assignment								
Received Own	-0.044***	-0.044***	-0.044***	-0.045***	-0.190***	-0.190***	-0.189***	-0.190***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.020)	(0.020)	(0.020)	(0.020)
Closest Peer Received		0.006				0.008		
		(0.005)				(0.014)		
No. Peers Received			-0.000				0.003	
			(0.001)				(0.002)	
Share of Peers Received				-0.019				0.032
				(0.014)				(0.055)
Mean Dep. Var.	0.089				0.225			
Std. Dev. Dep. Var.	0.105				0.287			

Notes: N= 2,060. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different linear regression. In columns (1)-(4), the dependent variable is the absolute value of the difference between the posterior belief about average peer salary and the actual average. *Received Own* is a dummy variable indicating whether the employee was randomly chosen to receive a signal about peer salary. *Closest Peer Received* is a dummy variable indicating whether the employee's closest peer (defined as the peer with whom the employee exchanges the most number of emails in Jan-Mar 2017) received a signal about peer salary before the employee started the survey. *No. Peers Received* denotes the number of employees in the employee's peer group who received a signal about peer salary before the employee started the survey. *Share of Peers Received* denotes the share of employees in the employee's peer group who received a signal about peer salary before the employee started the survey. The regressions control for the date of survey response, the total number of employees in the peer group and position title dummies. Columns (5)-(8) are equivalent to columns (1)-(4), but using manager salary instead of peer salary.

Table 3: Effects of Salary Perceptions on Behavior

	Effort and Performance			Career Moves			
	log($Hours$)	log($Emails$)	log($Sales$)	P(Quit)	P(Transfer)	log($Salary$)	$P(\Delta Title)$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post-Treatment:							
Log (Peer-Salary)$^{(i)}$	-0.943**	-0.431**	-0.731**	0.235**	0.093	0.004	0.114
	(0.472)	(0.210)	(0.297)	(0.107)	(0.106)	(0.052)	(0.123)
Log (Manager-Salary)$^{(ii)}$	0.150**	0.130***	0.106	-0.015	-0.003	0.002	0.012
	(0.074)	(0.041)	(0.122)	(0.022)	(0.030)	(0.011)	(0.029)
Pre-Treatment (Falsification):							
Log (Peer-Salary)	-0.205	-0.184	-0.191	-0.139	0.212	-0.001	-0.071*
	(0.542)	(0.289)	(0.412)	(0.218)	(0.163)	(0.005)	(0.040)
Log (Manager-Salary)	0.001	-0.101	0.063	-0.022	0.029	0.002**	0.009
	(0.114)	(0.071)	(0.160)	(0.050)	(0.029)	(0.001)	(0.010)
P-value H_0: (i)=(ii)	0.026	0.007	0.000	0.015	0.398	0.963	0.424
Cragg-Donald F-Stat.	29.8	204.0	98.2	203.7	203.4	203.6	203.3
Mean Outcome	5.98	35.57	0.48	0.05	0.09	0.92	0.10
Std. Dev. Outcome	1.88	44.93	0.23	0.21	0.28	0.70	0.30
Observations	602	2,060	791	2,060	2,060	2,060	2,060

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for two sets of instrumental variables regressions, following the specification described in section 2.2: in *Post-Treatment*, the dependent variable is the average behavior 90 days after the completion of the survey; in *Pre-Treatment (Falsification)*, the dependent variable is the average behavior before the completion of the survey. *Peer-Salary* is the posterior belief about the average peer salary, and *Manager-Salary* is the posterior belief about the average manager salary. The regressions control for three monthly lags of the dependent variable, (log) own salary and five productivity rating dummies. *Hours* is the daily number of hours worked. *Emails* is the daily number of emails sent. *Sales* is the sales performance index. $P(Quit)$, $P(Transfer)$ and $P(\Delta Title)$ are dummies for whether the employee leaves the firm, transfers inside the firm and changes position title, respectively. log($Salary$) is the logarithm of own salary at the end of the end (beginning) of the post-treatment (pre-treatment) period. For the dependent variables in logarithm, the mean and std. dev. reported in the bottom rows correspond to the pre-log values. Columns (1) and (3) have fewer observations because those dependent variables are only defined for a subsample.

Table 4: Sensitivity to Event Windows for Effects of Perceptions on Behavior

	Effort and Performance			Career Moves			
	$\log(Hours)$	$\log(Emails)$	$\log(Sales)$	P(Quit)	P(Transfer)	$\log(Salary)$	$P(\Delta Title)$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
3-Months Post-Treatment:							
Log (Peer-Salary)$^{(i)}$	-0.943**	-0.431**	-0.731**	0.232**	0.093	0.004	0.114
	(0.472)	(0.210)	(0.297)	(0.106)	(0.106)	(0.052)	(0.123)
Log (Manager-Salary)$^{(ii)}$	0.150**	0.130***	0.106	-0.016	-0.003	0.002	0.012
	(0.074)	(0.041)	(0.122)	(0.023)	(0.030)	(0.011)	(0.029)
6-Months Post-Treatment:							
Log (Peer-Salary)$^{(iii)}$	-1.200**	-0.437**	-1.064***	0.213*	-0.007	-0.212	0.120
	(0.591)	(0.209)	(0.347)	(0.127)	(0.157)	(0.176)	(0.123)
Log (Manager-Salary)$^{(iv)}$	0.112	0.073	0.233***	0.004	0.012	0.029	0.010
	(0.091)	(0.056)	(0.087)	(0.030)	(0.030)	(0.048)	(0.029)
P-value H_0: (i)=(iii)	0.748	0.407	0.400	0.603	0.740	0.577	0.980
P-value H_0: (ii)=(iv)	0.733	0.984	0.466	0.913	0.599	0.239	0.970
Mean Outcome	5.98	35.57	0.48	0.05	0.09	0.92	0.10
Std. Dev. Outcome	1.88	44.93	0.23	0.21	0.28	0.70	0.30
Observations	602	2,060	791	2,060	2,060	2,060	2,060

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. *3-Months Post-Treatment* shows the regression results from Table 3 – see its note for more details. *6-Months Post-Treatment* is identical except that the dependent variable corresponds to the behavior between the date of survey completion and 180 days later (instead of 90 days later).

Table 5: Effects of Salary Perceptions on Survey Outcomes

	Satisfaction		Redist. Pref.	Log(E[Future Salary])		Rank(Prod.)
	w/Pay	w/Job		+1 year	+5 years	
	(1)	(2)	(3)	(4)	(5)	(6)
Log (Peer-Salary)$^{(i)}$	-0.762*	-0.444	0.373*	0.071	0.280	0.044
	(0.433)	(0.491)	(0.216)	(0.090)	(0.176)	(0.040)
Log (Manager-Salary)$^{(ii)}$	-0.015	-0.086	0.008	0.025	0.166***	0.000
	(0.125)	(0.102)	(0.075)	(0.025)	(0.055)	(0.015)
P-Value (i)=(ii)	0.084	0.433	0.135	0.595	0.532	0.280
Cragg-Donald F-Stat.	253.6	254.3	254.3	253.5	255.3	250.5
Mean Dep. Var.	2.79	3.60	2.20	2.58	3.22	0.47
Std. Dev. Dep. Var.	0.92	0.78	0.57	0.51	0.59	0.22
Observations	2,030	2,027	2,027	2,033	2,026	1,999

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for a different instrumental variables regressions, following the specification described in section 2.2. The dependent variables are responses to survey questions elicited after the informational treatment. *Peer-Salary* is the posterior belief about the average peer salary, and *Manager-Salary* is the posterior belief about manager salary. All the dependent variables correspond to survey questions asked after the elicitation of the posterior beliefs. *Satisfaction with Pay* and *Satisfaction with Job* are measures in a 5-point scale from very dissatisfied (1) to very satisfied (5). *Redist. Pref.* measures preferences for within-firm redistribution in a 3-point scale from 1 (less) to 3 (more). *E[Future Salary]* corresponds to the expected salary 1 and 5 years in the future. *Rank(Prod.)* denotes the individual self-perceived position in the distribution of performance ratings in the firm. All regressions include the following control variables: the log of own salary, log of tenure, and sets of dummies for sales role, pay band, unit, productivity rating and position title.

Table 6: Effects of Perceived Manager Salary by Distance to Manager

	Effort and Performance			Log(E[Future Salary])	
	log($Hours$) (1)	log($Emails$) (2)	log($Sales$) (3)	+1 year (4)	+5 years (5)
Model 1 (Promo. Prob.):					
Log (Manager-Salary)					
Closer$^{(i)}$	0.212**	0.170***	0.437***	0.041	0.204***
	(0.099)	(0.052)	(0.154)	(0.030)	(0.059)
Farther$^{(ii)}$	-0.074	0.019	0.468	-0.008	0.086
	(0.093)	(0.104)	(0.755)	(0.033)	(0.092)
Model 2 (No. of Promo.):					
Log (Manager-Salary)					
Closer$^{(iii)}$	0.431*	0.185***	0.437**	0.008	0.200***
	(0.226)	(0.061)	(0.200)	(0.036)	(0.059)
Farther$^{(iv)}$	-0.016	0.068	0.516	0.057	0.134
	(0.135)	(0.062)	(0.526)	(0.038)	(0.096)
P-value H_0 : (i)=(ii)	0.040	0.243	0.972	0.216	0.229
P-value H_0 : (iii)=(iv)	0.170	0.212	0.908	0.322	0.560
Observations	602	2,060	755	2,033	2,026

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. In each model, *Log(Manager-Salary)* is interacted with dummies *Closer* and *Farther*, which denote the perceived distance between the own position and the managerial position. In Model 1, *Closer* indicates a probability of reaching the managerial position above 40%. In Model 2, *Closer* indicates managerial positions that are less than 5 promotions ahead. In columns (1)–(3), the dependent variables are the average behavior in the 90 days after the completion of the survey: *Hours* is the daily number of hours worked, *Emails* is the daily number of emails sent and *Sales* is the monthly sales performance – for more details about the specification, see the note to Table 3. In columns (4)–(5), the dependent variables are survey questions elicited after the posterior beliefs of manager salary: *E[Future Salary]* corresponds to the expected salary 1 and 5 years in the future – for more details about the specification, see the note to Table 5.