

How Much Does Your Boss Make?

The Effects of Salary Comparisons

Zoë Cullen

Harvard Business School

Ricardo Perez-Truglia

University of California, Berkeley

This Draft: June 14, 2021.

Abstract

The vast majority of the pay inequality in an organization comes from differences in pay between employees and their bosses. But are employees aware of these pay disparities? Are employees demotivated by this inequality? To address these questions, we conducted a natural field experiment with a sample of 2,060 employees from a multibillion-dollar corporation in Southeast Asia. We make use of the firm's administrative records alongside survey data and information-provision experiments. First, we document large misperceptions among employees about the salaries of their managers and smaller but still significant misperceptions of the salaries of their peers. Second, we show that these perceptions have a significant causal effect on the employees' own behavior. When they find out that their managers earn more than they thought, employees work harder, on average. In contrast, employees do not work as hard when they find out that their peers earn more. We provide suggestive evidence of the underlying causal mechanisms, such as career concerns and social preferences. We conclude by discussing the implications of pay inequality and pay transparency.

JEL Classification: D63, D83, J31, M12, M52, Z13.

Keywords: salary, inequality, managers, career concerns, transparency.

*Cullen: zcullen@hbs.edu, Rock Center 310, Boston, MA 02163. Perez-Truglia: ricardotruglia@berkeley.edu, 545 Student Services Building #1900, Berkeley, CA 94720. We thank the Editor and the anonymous referees for their very useful feedback. We are thankful for excellent comments from several colleagues and seminar discussants at MIT, Harvard, Stanford, Berkeley, Columbia, Northwestern, Yale, University of Chicago, UCSD-Rady, UCLA-Anderson, Berkeley-Haas, Harvard Business School, Wharton, Caltech, Brown, Einaudi Institute, Cornell, Dartmouth, Microsoft Research, Goethe, LMU, Heidelberg, Boston University, Luxembourg School of Finance, Paris School of Economics, Colegio de Mexico, Universidad de San Andres, University of Toronto, University of Zurich, Sao Paulo School of Economics, UCEMA, NES Political Economy Workshop, RIDGE Public Economics Workshop, Barcelona Summer Forum, Bogota Experimental Economics Conference, BEAM, Advances with Field Experiments 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).

1 Introduction

The vast majority of the pay inequality within organizations comes from salary disparities between employees and their bosses. We often hear about such pay disparities from the media, politicians, and policy makers. It is unclear, however, whether employees are aware of the pay inequality in their own firms. Pay secrecy policies, for example, may prevent employees from finding out how much their managers are making. Even if employees knew that their managers earned substantially more, it is unclear how that knowledge would affect employees' behavior. On the one hand, discovering that managers earn a lot may demoralize employees as a result of social preferences; they may feel jealous or resentful, or perhaps feel that the pay disparity is unfair. On the other hand, employees may find well-paid managers to be a source of motivation, by providing an extra incentive to work harder in hopes of being promoted to a managerial position with its large reward.

Are employees aware of how much their bosses get paid? And if employees believe their managers are paid handsomely, does it inspire them to greater effort or does it sap their motivation? In this study, we address these questions using a large-scale, high-stakes, natural field experiment in collaboration with a multibillion-dollar corporation. We study these questions using a research design that combines administrative data, incentivized surveys, and information-acquisition and information-provision experiments.

We designed a survey that elicits the respondents' perceptions about the average salaries of their managers and peers. For example, a junior analyst could be asked about the average salary of senior analysts (i.e., the managers) and about the average salary among the other junior analysts (i.e., the peers). These elicitations are incentivized, to ensure that it is in the best interest of participants to give thoughtful, honest answers. To measure the respondents' misperceptions about salaries, we can compare the respondents' perceptions to the true average salaries from the firm's administrative records. Our research design also sheds light on the nature of these misperceptions. To assess whether employees are aware of their misperceptions, we measure their level of certainty. And to assess whether they even care about others' salaries, we measure their willingness to pay for information, using an information-acquisition experiment (Becker et al., 1964).

Central to our research design is the information-provision experiment. After eliciting the prior beliefs, we randomized whether the employee would receive a signal about salary. For example, after eliciting beliefs of the manager salary, we randomized half of the subjects to receive a signal about the true manager salary. The signal consisted of the average salary in a random subset of employees in the position in question, which is different depending on the sample we draw. After the information-provision stage, we re-elicited beliefs regardless

of whether the subject received the signal or not. The first goal of the information-provision experiment was to measure how employees learn from information about salaries and whether they share that information with their coworkers. The second goal of the experiment was to create exogenous shocks to salary perceptions that allow us to estimate the causal effect of those perceptions on the employees' behavior as measured by administrative records. This allows us to test whether, consistent with the career concerns channel, a higher perceived manager salary translated into a higher effort or, consistent with the social preferences channel, demotivated employees.

Moreover, our research design allows us to provide some suggestive evidence on these specific causal mechanisms. We included a series of additional questions at the end of survey, after the information-provision experiments, to be used as survey outcomes. To assess the role of career concerns, we elicit expectations about the respondent's own future salary, using an incentive-compatible method. To examine the role of social preferences, we included questions related to employee morale (job satisfaction and pay satisfaction) and attitudes towards pay inequality. Our research design allows for an additional test of the career concerns channel. According to this channel, the effects of manager salary should depend on whether the employee can aspire to be promoted to that position or not. To test this hypothesis, we selected different managerial positions for different subjects: for example, we ask some junior analysts in investment banking about the average salary of senior analysts (a few promotions away), while we ask other junior analysts about the average salary of the chief economist (a higher number of promotions away). And to complement this data, we elicited the respondent's own perceptions of the number of promotions they would need to attain the managerial position, and the likelihood that they will be promoted to that position in the next five years.

We conducted the natural field experiment (Harrison and List, 2004) with a sample of 2,060 employees from a large commercial bank (referred to hereinafter as *the firm*) located in lower-middle-income country from Southeast Asia. The firm is comparable to other firms around the world in two key respects: pay transparency and pay inequality. The firm discloses limited information about pay and, according to our data, employees at this firm rarely discuss salaries with their coworkers. There is abundant evidence that both facts are true in most firms across the world, including but not limited to the United States (Trachtman, 1999; Edwards, 2005; Hegewisch et al., 2011; Glassdoor, 2016; PayScale, 2018). The degree of pay inequality inside this firm is not atypical either. For instance, the ratio between the 10th and 90th percentile of salaries is 0.21 in the firm compared to 0.19 for the average medium-sized firm in the United States (Song et al., 2019). And, as in most firms, the vast majority of the pay inequality is due to the salary differences between bosses and their subordinates

(Baker et al., 1988). For instance, the average manager salary that we elicited in our survey typically (i.e., 90% of the time) ranges between +114% and +634% of the employee’s own salary. In comparison, the peer average salary typically ranges between -16% and +16% of the employee’s own salary.

In the first set of results, we show that employees do have large misperceptions of the salaries of their managers and provide some evidence on the sources of these information frictions. We find that employees’ guesses about the average manager salary have a mean absolute error of 28%. Moreover, there is a systematic bias: on average, employees underestimate the manager salary by 14.1%. As a benchmark, the misperceptions of peer salary are still significant, but not as large: the mean absolute error is 11.5% and there is no systematic bias. We show that employees are aware of their own misperceptions, and that the misperceptions are not due to a lack of interest: some employees are willing to pay days’ and even weeks’ worth of salary for information about the salaries of their managers or their peers.

The results from this information-provision experiment illustrate that employees do learn from accurate information when it is provided to them. A simple Bayesian learning model indicates that employees put a 69% weight on the signal about manager salary we provided and only 31% on their prior beliefs. This finding indicates that the misperceptions must be due to a lack of access to information. Moreover, the results from the information-provision experiment identify a reason why employees do not have information in the first place: the lack of social learning. We rolled out the survey gradually over the course of two months. The staggered nature of the survey allows us to measure not only whether the information provided to an employee affects the employee’s own salary perceptions, but also whether that information diffuses to other employees connected to the participants who received information. We find no evidence of information diffusion: the information provided to an employee does not travel to their peers, not even to their closest peers.

In the second set of results, we show that the salary perceptions have a significant effect on behavior and provide evidence on the underlying mechanisms at play. To estimate the causal effects of salary perceptions on the employees’ own behavior, we leverage the exogenous variation in perceptions induced by the information experiment. The close collaboration with the firm, along with its rich administrative data, allows us to measure the effects of the information shocks on a number of different forms of behavior. We obtained data on two forms of effort: the number of hours an employee spent in the office (based on security data on all of the swipes in and out of the building) and the number of emails sent by the employees (based on data from the email servers). We also acquired data on one measure of performance: sales (for employees who have sales roles).

We estimate the effects of salary perceptions using a simple Instrumental Variables (IV)

estimator. To illustrate the intuition behind this estimator, consider the case of two respondents who, at the start of the survey, both underestimate the average manager salary by 20%. By chance, one of those respondents then receives a highly accurate signal about the average manager salary, while the other one receives no information at all. Based on the observed rates of updating, the employee who received no information continues to underestimate the average manager salary by 20%, while the employee who did receive the information underestimates the manager salary less – by 10%, let’s say. As a result, the information treatment amounts to a positive shock of 10% to the perceived manager salary. We can then measure how this 10% shock to that employee’s perceived manager salary affects her subsequent behavior. The IV estimator simply extends this logic to all of the respondents, not only the ones who underestimate the average salary by 20%.

We refer to comparisons between employees’ own salaries and the salaries of their managers as vertical comparisons, and to comparisons between employees’ own salaries and the salaries of their peers as horizontal comparisons. Regarding the vertical comparisons, the results from the IV regressions indicate that a higher perceived manager salary has a positive causal effect on the employee’s own effort and performance. We estimate that a 10% increase in perceived manager salary increases the average hours worked in the subsequent 90 days by 1.5%, implying a behavioral elasticity of 0.150 (p-value=0.042). The corresponding effects on the other measures of effort and performance are similar in magnitude: elasticities of 0.130 in the number of emails sent (p-value=0.001) and 0.106 (p-value= 0.383) in sales performance. The effects of the horizontal comparison, on the other hand, go in the opposite direction as the effects of the vertical comparison: a higher perceived peer salary has a negative causal effect on the employee’s own effort and performance. More precisely, a 10% increase in employees’ perception of their peers’ salaries has the number of hours they work by 9.4%, implying a behavioral elasticity of -0.94 (p-value= 0.045), with corresponding elasticities of -0.431 (p-value=0.041) in emails sent and -0.731 (p-value=0.014) in sales performance. Indeed, we can confidently reject the null hypothesis that the effects of horizontal and vertical comparisons are equal to each other: p-value=0.026 for hours worked, p-value=0.007 for emails sent, and p-value<0.001 for sales. All these results are robust to a number of checks, including but not limited to falsification regressions in an event-study fashion.

Next, we provide evidence about the mechanisms underlying the effects of the vertical comparisons. The fact that the perceived manager salary has a motivating effect suggests that the career concerns channel, which predicts positive effects, dominates over the social preferences channel, which predicts negative effects. Moreover, we provide more direct tests of these mechanisms using two different strategies.

The first strategy consists of measuring the effects of the salary information on survey

outcomes. We show that, as is consistent with the career concerns channel, when employees learn that their managers earn more, they become more optimistic about what their own salaries will be five years in the future. On the other hand, and contrary to what the social preferences channel would predict, we do not find any evidence that perceptions of manager salary have any effect on measures of employee morale (pay satisfaction and job satisfaction) or tolerance for pay inequality. In contrast, these same survey outcomes suggest that social preferences may be at play in peer comparisons: a higher perceived peer salary does have negative effects on employee morale (pay and job satisfaction) as well as on tolerance for pay inequality.

A second strategy exploits heterogeneity in the distance between the employee's own position and the managerial position. As is consistent with the career concerns channel, we find that the effects of the perceived manager salary are stronger for managerial positions that the employee can aspire to attain. When employees find out that managers who are a few promotions away earn more, they expect higher salaries in five years and they work harder. In contrast, when employees find about the high salaries of managerial positions they cannot aspire to attain, the effects are close to zero and statistically insignificant.

We discuss a number of implications for our findings. For example, our findings can help to explain why firms choose to load all of the incentives vertically (in the form of promotions) rather than horizontally (in the form of performance pay). Additionally, our findings have implications for the effects of disclosing salaries within a firm, as well as some recent policies around pay transparency. Moreover, we discuss some important caveats that one must keep in mind when extrapolating the findings from our specific setting to other settings. Indeed, the economic theories that we provide evidence for indicate that the results should change based on some mediating factors. Take for example the mechanism of career concerns. Economic theory suggests that this mechanism should be present only to the extent that there is opportunity for upward mobility in the organization. In other contexts, where employees have little expectation of reaching higher echelons in their organizations (e.g., Lyft drivers), disclosing the salaries of managers may not generate the same motivation as in our context, where movement along the career ladder is frequent.

This study is related to a recent but growing body of literature that looks at the effects of pay transparency and pay inequality. In a seminal study, [Card, Mas, Moretti, and Saez \(2012\)](#) conducted a field experiment with employees at the University of California in which they provide evidence that, consistent with models of social preferences among peers ([Frank, 1984](#); [Romer, 1984](#); [Lazear, 1989](#); [Akerlof and Yellen, 1990](#)), horizontal comparisons can demotivate employees: for workers who had salaries below the peer average, receiving a link to a website that publicized salaries decreased job satisfaction and the stated intention to

switch jobs. More recently, other studies have documented effects of pay inequality and pay transparency using natural experiments (Mas, 2017; Perez-Truglia, 2020; Dube, Giuliano, and Leonard, 2019), field experiments (Cohn, Fehr, Herrmann, and Schneider, 2014; Cullen and Pakzad-Hurson, 2016; Breza, Kaur, and Shamdasani, 2018), and laboratory experiments (Bracha, Gneezy, and Loewenstein, 2015; Huet-Vaughn, 2017).

Our study advances the existing literature in two ways. First, whereas the previous work focused on horizontal comparisons, our study investigates both vertical and horizontal comparisons. Horizontal pay inequality accounts for a small share of the overall inequality within firms (Baker et al., 1988). In the firm we studied, for example, less than 5% of the salary inequality is horizontal. This distinction between vertical and horizontal comparisons is important, as we find that the two types of comparisons have effects in opposite directions, through different mechanisms, and thus have different implications.

The second way we contribute to this literature is by providing direct evidence of salary misperceptions. 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), consisting of a survey of 326 managers from four privately owned U.S. companies. He finds those respondents systematically underestimate the salaries of those in higher positions but do not systematically underestimate the salaries of those in their same position. Another notable study, Kiatpongsan and Norton (2014), elicits perceptions about aggregate salary statistics: using survey data from 16 countries, they show that most households underestimate the pay ratio between chief executive officers and the average unskilled workers. We contribute to this literature by providing novel evidence on some questions that remain largely unexplored in economics, such as what employees know about salaries and how the information travels among coworkers. Our close collaboration with the firm allowed us to provide measurements that had proved elusive in previous work. We are able to match incentivized survey responses to administrative records, which allows us to measure salary misperceptions directly. Additionally, our unique experimental design enables us to disentangle the sources of the misperceptions by measuring employees' willingness to pay for information and the diffusion of salary information.

Our study also contributes to literature about career concerns and salary dynamics (Lazear and Rosen, 1981; Harris and Holmstrom, 1982; Rosen, 1986; Gibbons and Murphy, 1992; Holmstrom, 1999; Dewatripont et al., 1999; Gibbons and Waldman, 1999a,b). Although there is a rich body of theoretical work on this topic, there is little direct evidence about these mechanisms. We contribute by testing some of the central predictions of these models. We show that employees form their expectations about future salaries based on what they think their bosses are making; and that employees work harder when they find

out that a position to which they aspire to get promoted offers higher salaries. More broadly, our study also relates to a literature on the determinants of employee morale (Dellavigna, List, Malmendier, and Rao, 2019) and on the effects of relative income on job satisfaction (Godechot and Senik, 2015; Clark, Frijters, and Shields, 2008) and happiness (Senik, 2004).

The rest of the paper proceeds as follows. Section 2 describes the research design, including the institutional context and data sources. Section 3 presents the results on salary misperceptions. Section 4 discusses the effects of perceived salaries on behavior. Section 5 discusses the generalizability of the results and its implications for pay transparency. The last section concludes.

2 Research Design

In this section, we discuss the most important aspects of the survey design.

2.1 Conceptual Framework

Why would employees want to be informed about the salaries of their managers and peers? How would that information affect their effort? In this section, we provide a toy model that illustrates two mechanisms that motivated the research design: career concerns and social preferences.

We introduce a two-period model in which the utility of the employee is a linear function of the salary in the first period (w_1) and the expected salary in the second period (w_2), with discount factor $\delta \in (0, 1]$. If promoted, the employee's expected salary is given by \tilde{w}^z . The probability that the employee is promoted depends on the employee's chosen effort ($e > 0$) and his or her ability parameter ($\theta > 0$): $e \cdot \theta$. We assume a quadratic effort cost: $c(e) = \frac{1}{2}e^2$.

Let \bar{w}^{mgr} denote the employee's perception about the average salary of managers. The employee's expected salary if promoted is an increasing function of perceptions about the average salary of managers: $\tilde{w}^z(\bar{w}^{mgr})$, with $\frac{\partial \tilde{w}^z}{\partial \bar{w}^{mgr}} > 0$.¹ This assumption is intended to capture that the employee updates about the wage-setting policy of the firm; also, the employee may plan on using the information strategically to negotiate his or her salary if promoted.

If the employee is not promoted, he or she can still renegotiate the salary, denoted \tilde{w}^r . Let \bar{w}^{peer} denote perceptions of the average salary of peers. We assume the expected raise if not promoted is an increasing function of the average salary of peers: $\tilde{w}^r(\bar{w}^{peer})$, with $\frac{\partial \tilde{w}^r}{\partial \bar{w}^{peer}} > 0$. This assumption is intended to capture that, if not promoted, the employee can try to use the information on peer salary strategically to negotiate a raise.

¹We assume this expected salary is always greater than the employee's current salary ($\tilde{w}^z > w_1$). Likewise, we assume that the renegotiated salary cannot be higher than the salary if promoted ($\tilde{w}^z > \tilde{w}^r$).

We model social preferences following [Dellavigna et al. \(2019\)](#) and [Breza et al. \(2018\)](#), by introducing a morale term in the utility function that depends on relative pay vertically ($\bar{w}^{mgr} - w_1$) and horizontally ($\bar{w}^{peer} - w_1$), with parameters $\psi^{mgr} > 0$ and $\psi^{peer} > 0$ that capture the intensity of social preferences.² Putting all the pieces together, the employee's expected utility is given by the following expression:

$$V(e) = w_1 - \frac{1}{2} \cdot e^2 + \delta \cdot (e \cdot \theta \cdot \tilde{w}^z + (1 - e \cdot \theta) \cdot \tilde{w}^r) - \psi^{mgr} \cdot e \cdot (\bar{w}^{mgr} - w_1) - \psi^{peer} \cdot e \cdot (\bar{w}^{peer} - w_1) \quad (1)$$

Employees choose effort to maximize expected utility given by equation (1). Our functional form assumptions ensure the objective function is concave, $U_{ee} < 0$, and thus the interior solution to the maximization problem is a local maximum. Solving for the first order conditions yields the following expression of optimal effort (e^*):

$$e^* = \delta \cdot \theta \cdot (\tilde{w}^z(\bar{w}^{mgr}) - \tilde{w}^r(\bar{w}^{peer})) - \psi^p \cdot (\bar{w}^{peer} - w_1) - \psi^{mgr} \cdot (\bar{w}^{mgr} - w_1) \quad (2)$$

Next, we provide some simple comparative statics on the effects of salary information on effort. We start with the effects of information on manager salary:

$$\frac{\partial e^*}{\partial w^{mgr}} = \underbrace{\delta \cdot \theta \cdot \frac{\partial \tilde{w}^z}{\partial w^{mgr}}}_{\text{Career Concerns}} \quad \underbrace{-\psi^{mgr}}_{\text{Social Preferences}} \quad (3)$$

The first term, which is always positive, captures the career concerns channel. The intuition is straightforward: the employee wants to work harder as a response to the higher expected rewards from being promoted. The second term, which captures the social preferences channel, is always negative: the vertical inequality demoralizes the employee and hence reduces effort.

Next, we discuss the comparative statics for the information on peer salary:

$$\frac{\partial e^*}{\partial w^{peer}} = \underbrace{-\delta \cdot \theta \cdot \frac{\partial \tilde{w}^r}{\partial w^{peer}}}_{\text{Career Concerns}} \quad \underbrace{-\psi^{peer}}_{\text{Social Preferences}} \quad (4)$$

The career concerns channel is reflected in the first term, which is always negative. The intuition is straightforward: since the employee knows that he or she can still get a raise even if not promoted, the incentive to work hard diminishes. The social preferences channel is captured by the second term, which is always negative, implying that the horizontal inequality demoralizes the employee and hence reduces effort.

In the rest of the paper, we will refer to career concerns and social preferences as two

²There are a number of models in which employee morale increases with relative pay ([Frank, 1984](#); [Romer, 1984](#); [Lazear, 1989](#); [Akerlof and Yellen, 1990](#)). For example, employees may follow a reciprocity norm ([Akerlof, 1982](#); [Gneezy and List, 2006](#)) and thus feel obligated to work harder if they are paid more than their peers and less obligated to work hard if they are paid relatively worse.

distinct channels. On the one hand, you should keep in mind that this is a toy model and thus not intended to capture all the different mechanisms that could be at play. For instance, an alternative instrumental reason why employees may care about the salary of their peers is that they use it to infer their market value and thus to decide whether to seek or accept outside offers.³ On the other hand, the toy model makes a number of additional predictions that we can test with our experiment and thus provide evidence in favor or against these specific mechanisms. First, the career concerns channel predicts that a higher manager salary should increase expected own future salary.⁴ Second, the career concerns channel predicts that the effects of manager salary should be mediated by whether the employee can aspire to be promoted to the managerial position he or she is learning about.

2.2 Institutional Context

To design and conduct the experiment, we collaborated with a private commercial bank from a lower-middle-income country in Southeast Asia. To keep the identity of the firm secret, we refrain from being specific about the firm’s characteristics. This firm has millions of customers, billions of dollars in assets and in revenues, and thousands of employees. These employees are based in two headquarters and in hundreds of branches dispersed over the country.

This firm is comparable to other large firms around the world in some relevant respects. Regarding pay inequality, the ratio between the 10th and 90th percentile of salaries is 0.21 in this firm, whereas it is 0.19 for the average medium-sized U.S. firm (Song et al., 2019). The inequality in this firm is also typical in that only a small part of it is horizontal. A simple inequality decomposition suggests that less than 5% of the pay inequality is horizontal, which is in the same order of magnitude as in other organizations studied in the literature (Card et al., 2012; Baker et al., 1994).⁵ Indeed, for a more direct comparison, we can study the pay differences between employees and their managers for specific positions. For instance, the ratio between the salary of a senior relationship manager and their subordinate, a personal retail banker, was 1.5 in this firm. According to 2017 data from Glassdoor, the corresponding ratio was also 1.5 for U.S. employees at Bank of America.

Regarding pay transparency, the firm does not have an open salary policy.⁶ Several

³In the spirit of Jovanovic (1979), if employees use the peer salary information to infer their outside value, bad match employees are more likely to leave when they learn their peers are paid better. Indeed, Caldwell and Harmon (2018) provides some evidence in support of this mechanism.

⁴There are two reasons for that result. The direct effect is that salary expectations conditional on promotion rise with manager salary. The indirect reason is that the effort rises endogenously with manager salary, thus increasing the chances of achieving a promotion.

⁵For details, see Appendix C.1.

⁶The firm discloses some information about pay, but this information is too vague to form a decent guess

surveys corroborate this pattern of pay secrecy around the world. For example, a 2003 survey of Fortune-1,000 firms shows that only 3.5% of the surveyed firms had open salary policies (Lawler, 2003); a survey of about 1,000 companies indicates that only 3% have open salary policies and less than a quarter disclose data on salary ranges (Scott, 2003). A survey of 7,100 managers from the United States and other countries indicates that only 6% have open salary policies (PayScale, 2018). Indeed, the standard employment contract at this firm explicitly prohibits employees from sharing salary information. Many organizations around the world have similar policies (Day, 2007). In a survey of private sector employees from the United States, for example, more than 60% report that their employer discourages or prohibits employees from discussing salaries with coworkers (Hegewisch et al., 2011). According to our survey data, 45% of employees never discuss salaries with coworkers.⁷ Similar patterns have been documented around the world. For instance, a survey of 1,022 employees from the United Kingdom found that 48% discuss salaries with their peers (Burchell and Yagil, 1997); and a 2017 survey of Americans aged 18-36 years show that 70% report to never discuss salaries with coworkers (Gee, 2017).

Among the observable characteristics, there are some features of this firm that may be less representative. For example, the majority of the employees at this firm are female.⁸ The firm may be different in some unobservable characteristics too. We return to this topic in Section 5.1, where we discuss how the results from this specific firm should be extrapolated to other settings.

2.3 Survey Design: Managers and Peers

The survey revolves around the average salaries of two groups: managers and peers. To identify a managerial position for each employee, we used multiple sources of administrative data. The criteria for who we considered a manager can be summarized as follows: 1) The managerial position had to be occupied by someone in the respondent's unit; 2) The managerial position had to be higher than the respondent's position; 3) The managerial position had to have an oversight role over the respondent, such as conducting performance evaluations or approving leaves of absence.

about the average salaries of peers and managers. For instance, the firm discloses the existence of a 10-point pay band scale, but the minimum and maximum salaries in these boundaries are not disclosed, and they even overlap quite a bit.

⁷More 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. Since this type of behavior is frowned upon by the employer, it is probably under-reported in surveys.

⁸At the time of the experiment, the share of female employees at this firm was 71%. In comparison, the U.S. workforce of Bank of America was 53% female. Source: <https://about.bankofamerica.com/en-us/what-guides-us/our-global-workforce.html>.

According to the career concerns mechanism from Section 2.1, the effects of manager salary should be present only to the extent that the employee can aspire to be promoted to the managerial position. To test this hypothesis, we pick a different managerial position for different respondents. For example, a junior analyst in investment banking could be asked about the average salary of senior analysts (a position that is a few promotions above them) or about the salary of the chief economist (a higher number of promotions away). And to complement this data, we included two questions in the survey to elicit the respondents’ own perceptions of the distance between their own positions and the managerial position: the number of promotions needed to attain the managerial position and the likelihood of being promoted to that position within five years.

We use a definition of peer group that is close to the definition used in other studies (Card et al., 2012; Cohn et al., 2014; Dube et al., 2019; Cullen and Pakzad-Hurson, 2016; Breza et al., 2018): employees with the same position title who work in the same unit. The peers of a junior analyst in investment banking, for example, would be the other junior analysts in investment banking. Because they have the same title, these employees should have the same powers and the same responsibilities. Employees typically work in close physical proximity to their peers and on some occasions they may even need to collaborate with each other. In the survey, we provide specific instructions about the definition of each group. In the case of peer group, for example, we state the full position title, the full name of the unit, and the number of employees currently working in that peer group.

2.4 Survey Design: Training

A sample of the full online survey is included in Appendix A.⁹ The first module of the survey was designed to teach the subjects some basic concepts that would be useful for the rest of the survey. It begins with an explanation of how the incentivized questions work. All of the accuracy rewards in the survey were set up using the traditional quadratic loss function calibrated to award up to \$2.61 per question – this amount, as well as all other monetary amounts discussed in the paper, have been transformed to United States dollars using PPP-adjusted exchange rates from April 2017.¹⁰ To make sure that subjects understood the incentives for truth-telling, in the training portion of the survey we explained that it is in

⁹To protect the identity of the firm, we removed all identifying information from this survey instrument, including the formatting.

¹⁰More specifically, if x represents the accuracy (e.g., $x = 0.01$ if the subject’s guess is off by 1%), then the reward is given by the following expression: $\$2.61 \cdot (1 - \max\{\min\{10 \cdot x^2, 1\}, 0\})$. We did not disclose the formula for accuracy rewards explicitly – this approach is supported by recent evidence indicating that the highest degree of truth telling occurs when the technical details about the incentives are withheld (Danz et al., 2020).

their best interest to report truthfully and provided a practice question (about a topic not related to salaries) to cement this knowledge. Moreover, to make the incentives for truth-telling as salient as possible, we added a sentence reminding subjects about the accuracy rewards at the end of each incentivized question.

One potential concern with incentivized questions is that subjects may learn whether their guesses are right or wrong based on the accuracy rewards. We took several measures to shut down any significant possibility of learning. First, we did not provide feedback to the subjects on whether they got any specific question right or wrong – subjects do not even find out how much they got paid for any specific incentivized question. Moreover, we designed the rewards in a way that makes it almost impossible for a subject to infer anything about the accuracy of their guesses.¹¹

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 feature of compensation that is most salient and most important to employees. For instance, when a new employee joins the firm, the monthly base salary is the key figure written in the contract. Indeed, the base salary accounts for over 90% of the total compensation for the subjects in our sample.¹² To confirm that respondents understood the definition of salary, we also included an exercise where they need to guess their own salary.¹³ This question was also intended to convey that the surveyor already knows the salary of the respondent, thus undermining any inclination on the part of the respondents to misreport their salaries in order to avoid revealing them to the researchers.

¹¹We did not disclose the exact formula used for the accuracy rewards, this limiting the subject’s ability to infer anything from the rewards they receive. Moreover, we designed the survey rewards so that subjects cannot even infer what was the total amount earned through their accurate guesses. A few weeks after the survey, the participants receive a direct deposit in the full amount for the survey participation. The formula for the final reward is provided to subjects in the last page of the survey: the sum of the rewards earned in the survey for accuracy, plus a fixed fee of \$6.52 and a surprise amount picked at random, following a uniform distribution in the range \$3.26-\$14.35.

¹²Using transaction data provided by the bank for the employees who participated in our experiment, we find that the base salary as a fraction of total compensation has a median of 97.2% and an average of 93.3%. The second largest source of compensation for employees who have some form of sales role is sales commissions, but they tend to be small relative to the base salary. Other forms of performance pay can be substantial for employees in the highest positions, such as C-suite bonuses, but those employees were excluded from participation in our study.

¹³We asked respondents to guess their own salaries for the month of March 2017 and offered a reward for accuracy. On the next screen, we showed the participant’s guess as well as the true salary. If the respondent’s guess was not within 5% of the true salary, we showed them an additional screen re-explaining the definition.

2.5 Survey Design: Salary Perceptions

The two main modules, on manager salary and peer salary, follow the structure below:

- **Step 1 (Elicit Prior Belief):** We asked respondents about the average monthly base salary among peers/managers. To elicit truthful responses, we offered a reward for accuracy. To get a sense of how certain respondents felt about their guesses, we also 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 acquire the following piece of information: the average salary over a random sample of five managers/peers. To elicit this information in an incentive-compatible way, we employed the multiple price list variation of the Becker-DeGroot-Marschak (BDM) method (Becker et al., 1964). This method consists of having respondents face five hypothetical trade-offs between information, and a reward amount that varies. For each trade-off, the respondent can choose to either see the piece of information or add an amount of money to their survey rewards (i.e., the “price” of the information). Since all employees must have accounts in the bank where they work, the monetary rewards could be deposited directly into a respondent’s bank account. The five trade-offs effectively capture different pricing for the information: \$1.3, \$6.5, \$26.1, \$130.5, and \$652.3.¹⁴ We explained to subjects that making truthful choices was in their best interest because there was a small probability that one of the five trade-offs would be randomly selected to be executed. For the 1% of respondents who had their choice executed (for the manager salary or the peer salary), the survey was automatically terminated; thus, they are excluded from the subject pool. The other 99% of respondents continued with the rest of the survey.
- **Step 3 (Information-Provision Experiment):** For each subject, we calculated the two signals described in the previous step: the average salary over a random sample of five managers, and the average salary over a random sample of five peers. We then cross-randomized whether the subject would get to see each of the signals. Each subject faces a 50% probability of seeing each signal. To avoid respondents 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 employees participating in this survey would be randomly chosen to receive the signal about manager/peer salary for

¹⁴We calibrated this scale using a small pilot survey that elicited willingness to pay for information with an open-ended and non-incentivized question.

free. In the following screen, we let the subjects know whether they were chosen to receive the signal or not.

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

The module for peer salary appears first and the module for manager salary appears next. With respect to the information-provision experiment, we cross-randomized the two pieces of information, which resulted in four treatment groups: one group received a signal about the average salary of their peers but no salary information about their manager; one group received a signal about the salary of their manager but not those of their peers; one group received information about both their peers' and manager's salaries; and one group received no salary information.

Our survey elicited beliefs about the average salaries in those specific groups (managers and peers). In practice, employees may be interested in other moments of the distribution, such as the median, minimum, or maximum. This design choice was based on interviews with employees who were not invited to the survey and also managers from the human resources division, all of whom indicated that the information about averages was most relevant for them. 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. Employees may be interested in other types of salaries besides their managers and peers. Our interviews at the firm indicated that these are the types of positions that employees are most interested in. Indeed, we included a question in our survey eliciting the positions that the employees are most interested in learning about. The responses to this question are consistent with the anecdotal evidence.¹⁵

2.6 Behavioral Outcomes

The main goal is to estimate the effect of salary perceptions on the employee's effort and performance. We have two proxies for the effort of the employee. We collaborated with the different units of the organization to create three key measures. The first behavioral outcome is observed for employees who work in the headquarters (29% of the sample). Employees there

¹⁵When asked about the piece of salary information they would be most interested in learning about, roughly 50% of subjects ranked highest their own position, 45% ranked highest the managerial positions, and less than 5% ranked highest other positions.

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, minutes, and seconds that each employee spends at work on a daily basis. We then compute the average of hours worked per day.¹⁶

The second measure of effort is observed for every employee in the sample. We scraped the email servers of the company in real time, collecting data on the emails sent and received by all employees.¹⁷ Our measure of effort is defined as the total number of emails sent by the employee on a daily basis. The advantage of this measure over the alternative, of hours worked, is that it is available for the entire subject pool. While the number of emails may not be a great 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. Employees need to send emails to clients or coworkers for most of their duties, such as reaching out to new clients, or obtaining internal approvals for loans or credit cards. Last but not least, due to company policy, employees must leave an email trail for some of their tasks. For example, after calling a client to offer a product, employees are required to follow up with an email containing the information shared over the phone. Consistent with the above anecdotal evidence, the number of emails is positively and significantly associated with the alternative measure of effort, which is the number of hours spent in the office.¹⁸ Last, while to the best of our knowledge our study is one of the first in economics to use real-time email data to measure effort, there is a recent but growing literature in information science demonstrating that the email behavior predicts employee performance (for a literature review see [Wen et al., 2020](#)).

As discussed in [Dellavigna et al. \(2019\)](#), there are different margins of effort, and while some of those margins may be elastic to experimental interventions, some other margins may be quite inelastic. For example, in the experiment conducted by [Dellavigna et al. \(2019\)](#), subjects were much more elastic to incentives when deciding to stay on the job after hours. We believe our measures of effort – hours in the office and emails sent – are likely to be elastic, too. Employees seem to have quite a bit of discretion in how long to work (leaving a bit earlier or staying after hours), and in how many emails to send and respond to. The

¹⁶Since some employees work during weekends, we average over all seven days of the week. The average number of hours worked in our sample is in the ballpark for the typical office jobs. According to the U.S. Bureau of Labor Statistics, the average American works 44 hours per week ([CNBC, 2017](#)). Divided by seven days of the week, this comes down to a daily average of 6.29 hours, which is in the same order of magnitude as the corresponding average in our sample (5.98 hours, from column (1) of [Table 2](#)).

¹⁷Due to the sensitive nature of the data, we did not retrieve any information on the content of the emails.

¹⁸Results reported in [Appendix C.2](#).

hours worked and emails sent are not explicitly monitored by the company and employees are not rewarded or punished for them. As a result, these forms of effort may be more reactive to factors such as career concerns and social preferences.

We have one measure of performance, for employees who have a sales role (38% of the sample). The firm has detailed data on the sales revenue of each employee at the monthly level. We use the firm’s standard formula to aggregate sales across the different products offered by the firm (e.g., credit cards, loans, mortgages). Our measure of performance is defined as the employee’s rank in the monthly distribution of normalized sales revenues, ranging from 0 (lowest) to 1 (highest). As a validation exercise, we show that the number of emails is positively and significantly associated to this measure of sales performance.¹⁹

In addition to effort and performance, the information about manager and peer salary may affect career outcomes such as retention and raises. We use administrative data from HR to track those secondary outcomes. We began collecting all the administrative data three months before launching the survey. As a result, in addition to post-treatment outcomes, we can measure the pre-treatment outcomes, which can be used as control variables to improve precision as well as for falsification tests.

2.7 Survey Outcomes

In addition to the behavioral outcomes, we can measure the effects of the information on a series of outcomes elicited at the end of survey (i.e., after the information-provision experiments). These survey outcomes can provide evidence of the underlying mechanisms at work.

According to the career concerns channel from Section 2.1, if employees learn that their managers get paid more and expect to reach that same position eventually, then they should become more optimistic about their own future salaries. To test this hypothesis, we elicit expectations about respondents’ own future salaries one year and five years in the future, using an incentive-compatible method.²⁰

According to the social preferences channel from Section 2.1, the salary information may affect employee morale and tolerance for pay inequality. We included three questions meant to gauge this mechanism. Following prior work (Clark and Oswald, 1996; Card et al., 2012), we elicit the employee’s satisfaction as a proxy for employee morale. The first question is about satisfaction with pay specifically: “How satisfied are you with your current salary?”

¹⁹Results reported in Appendix C.2.

²⁰To incentivize truthful responses, it was not practical to offer rewards by comparing the guesses to the actual future salaries, because we would need to wait five years to calculate the rewards. Instead, we told respondents that we would compare their guesses to our own predictions of their future salaries (based on our administrative data and predictive models).

Responses to this question used a 5-point scale from very dissatisfied (1) to very satisfied (5). The second question, about overall job satisfaction, uses the same 5-point scale: “Taking all aspects of your job into account, how satisfied are you with your current job?” The third question, on tolerance for pay inequality in the firm, is an adaptation of a traditional question from the literature on preferences for redistribution (Cruces et al., 2013): “Across the thousands of [Bank Name] employees, salaries vary with the nature of the work, education, experience, responsibilities, etc. What do you think of wage differentials in the company today?” The possible answers were: (1) They are too large; (2) They are adequate; and (3) They are too small. Higher values of this outcome thus indicate higher tolerance for pay inequality.

2.8 Survey Implementation

To construct the subject pool, we started with the universe of employees, numbering in the thousands, and excluded some employees because of data limitations or by request of the firm.²¹ After these filters, we were left with a sample of 3,841 employees, all of whom were invited via email to participate in our survey. A sample of the invitation email (stripped of formatting and identifying information) is presented in Appendix B. The invitation email stated that the survey typically takes less than 30 minutes, and that survey participants would receive, on average, \$30 as rewards for participating in the survey, but did not include any specific information about the content of the survey. Participation was not compulsory, but employees were strongly encouraged to participate.²² The email invitations were sent out gradually over the course of two months, and we collected survey responses from the first week of April 2017 to the first week of June 2017. This staggered timing was designed to measure the diffusion of the salary information. Of the 3,841 invitations sent out, 2,060 employees completed the main module of the survey, corresponding to a 53.6% response rate.²³

²¹The firm asked us to exclude employees from the highest pay bands (i.e., the highest executives), employees who had joined the firm in the past 6 months, and employees from smaller divisions. Further, we exclude a few employees due to data limitations: a small minority of employees who belonged to small peer groups, and a minority of employees who could not be matched to managerial positions.

²²The invitation email listed the endorsements of three of the firm’s high-level executives and, additionally, the heads of each unit reached out in separate emails to encourage participation.

²³This sample already excludes the subjects who were randomly assigned to have their choices in the information-shopping scenarios executed – for more details, see Appendix C.3.

2.9 Descriptive Statistics and Randomization Balance

The subject pool comprises employees from different pay bands, with all types of roles (e.g., analysts, technicians, tellers, sellers, clerks, receptionists). On average, subjects are 29 years old and have been working at the firm for five years. 73% of them are female and 86% have a college or higher degree. The median number of employees in a peer group is 19, and the 25th and 75th percentiles are 10 and 32. Appendix C.3 provide further descriptive analysis. For instance, we show that, consistent with successful random assignment, there is balance in observables across treatment groups. And we show that the subject pool is representative of the universe of employees in all of these observable characteristics.

There are some additional features of the setting that are particularly important to keep in mind. First, while the salary differences between peers tends to be small, there are large differences between the salaries of the employees and their managers. For instance, the mean absolute difference between the subjects' own salaries and the average peer salary is 11.7% of the subjects' own salaries; in comparison, the mean absolute difference between the subjects' own salaries and the average manager salary is 315% of their own salaries.²⁴ Second, there are plenty of opportunities for upward mobility (the annual promotion rate is around 16.5%) and employees seem to be well aware of this fact.

3 Results: Beliefs about Manager and Peer Salary

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

3.1 Accuracy of Prior Beliefs

We measure misperceptions by comparing the employees' salary guesses against the true figures from the administrative records of the firm. Figure 1.a shows misperceptions of average manager salary. Only a small share (12%) of respondents guess the average manager salary within $\pm 5\%$ of the truth. The rest of the respondents miss the mark, often by a large margin: the mean absolute error is 28%. As a benchmark, Figure 1.b shows misperceptions of average peer salary. While still significant, the misperceptions of peer salary are smaller than the misperceptions of manager salary. The fraction of employees who can guess their peers' salaries within 5% (32% of respondents) is 2.6 times the fraction of employees who can guess the managers' salaries (12% of respondents). The mean absolute error for manager

²⁴This is just the tip of the iceberg, as we did not ask any employees about managerial positions that were well above their pay grade (e.g., asking tellers about the pay of the CEO).

salary (28%) is 2.4 times the mean absolute error for peer salary (11.5%). Moreover, while there is a systematic negative bias of 14.1% in perceptions of manager salary, there is an average overestimation of peer salary but of only 2.5% (p-value<0.01).

We can provide some robustness checks too. We find that the vast majority of employees were able to guess their own salaries almost exactly on their first try, thus confirming the anecdotal evidence that base salary is a salient aspect of compensation in this firm.²⁵ We also measured misperceptions in a follow-up study with a different sample and somewhat different methods. Despite the differences, the misperceptions are in the same order of magnitude.²⁶ We also show that the misperceptions are not driven by any specific subgroup, such as female employees, employees at the bottom of the hierarchy, or employees in large peer groups.²⁷ Last, as shown in Figure 1.c, the misperceptions of manager and peer salary are largely unrelated to each other: their correlation coefficient is statistically significant (p-value=0.007) but small in magnitude (0.059).

We show that employees are largely aware of their misperceptions. Our survey elicited the probability that the true salaries fall within certain bins around the respondent's guess. We find that employees are largely aware that they do not know the manager and peer salary perfectly. For example, respondents on average think that there is a 32.2% probability that the true manager salary falls within 2.5% (i.e., +/-2.5%) of their guesses, while the corresponding probability for peer salary is 33.8%. In reality, the probability of guessing the manager salary within 2.5% of the truth is only 8% of the guesses. For peer salary, the fraction of guesses that are that accurate is only 16.1%. We interpret this as evidence that, despite being aware that they are far from perfectly accurate, employees are overconfident about their accuracy.

Our favorite interpretation of the large misperceptions reported above is that employees have little information about salaries beyond knowing their own salaries. In the case of perceptions of peer salary, the subjects' own salaries are significantly informative about the average peer salary. Indeed, we find that a significant fraction of respondents seem to be reporting their own salary and, moreover, for those who do not, they are not doing any better than if they had reported their own salary.²⁸ While reporting one's own salary is a

²⁵Results reported in Appendix C.4.

²⁶Cullen and Perez-Truglia (2018) provided different incentives (rewarding responses within 5% of the truth instead of using the quadratic scoring rule), used larger reward amounts (earning up to \$63, instead of up to \$2.61), and elicited a different belief (the average salary among a specific sample of five peers, instead of the average among all peers). Despite the differences, the results are robust: the mean absolute error for peer salary is 11.5% in this survey versus 14.6% in Cullen and Perez-Truglia (2018).

²⁷Results reported in Appendix C.5.

²⁸We find that 35% of subjects report a guess for average peer salary within 5% of their own salaries. If all employees had reported their own salaries as a guess for the average peer salary, the mean absolute error would have been 11.4% (vs. 11.5% in reality), and the bias would have been -0.4% (vs. 2.5% in reality).

reasonable idea when guessing the peer salary, one’s own salary would be a poor guess for the average manager salary, which could explain why employees fare so much worse at guessing manager salary.

3.2 Willingness to Pay for Salary Information

The misperceptions of the salaries of peers and managers may simply reflect a lack of interest on the information. The data on willingness to pay for information can shed light on this question.

For each piece of information (manager and peer salary), respondents faced five different trade-offs between information and a reward, effectively setting the price of the information. We find that the vast majority of subjects (80% for manager information and 85% for peer information) made selections that are consistent across scenarios. Following the standard practices, the following results focus on subjects with consistent responses.²⁹ Figure 1.d shows the distribution of the willingness to pay for the signal about manager salary. The results suggest that while some employees see little value in information about manager salary, a substantial fraction of them value it a lot. On the one extreme, 22.9% of employees are not willing to pay more than \$6.5 for the signal about manager salary, an amount that is typically less than an hour’s worth of salary.³⁰ On the other extreme, 24.9% of employees are willing to pay more than \$652 for the information, which for most employees constitutes more than a week’s worth of salary. This substantial willingness to pay for salary information suggests that a great deal of the misperception arises because acquiring information is difficult rather than because employees are uninterested. Figure 1.e shows the corresponding results for the willingness to pay for peer information instead of manager information. On aggregate, the willingness to pay for manager and peer information seems to be quite similar. However, there are substantial discrepancies at the individual level: as shown in Figure 1.f, there are plenty of employees who are interested in the manager information but not the peer information, and vice versa.³¹

²⁹Responses are consistent if they are monotonic in the price list (Andersen et al., 2006). For example, if the respondent is willing to forgo \$130.46 for the information, he or she should be also willing to forgo \$26.09 for the manager information. The rates of consistent responses are in the same order of magnitude as the corresponding rates reported in other studies employing the price-list method (Andersen et al., 2006; Allcott and Kessler, 2019; Fuster et al., 2018).

³⁰We refrain from providing more precise information to avoid revealing sensitive information about compensation levels at the firm. Also, 19.3% of the respondents were not willing to buy the information even for the lowest price of \$1.3. In principle, some of these subjects could have a negative willingness to pay: that is, they would like to pay to avoid seeing the information. We did not extend the price list to include negative prices because this information aversion seemed largely inconsistent with what we saw in qualitative interviews with employees.

³¹More precisely, the correlation coefficient between the willingness to pay for manager and peer informa-

The average willingness to pay for the information is quite substantial. Since we elicit the willingness to pay in intervals, we must make a choice about how to compute the average. The most conservative approach consists of using the lower bound of the interval for each individual. That conservative estimate puts the average valuation at \$191 for the manager information and \$197 for the peer information. If we use less conservative approaches, we find somewhat larger magnitudes but still in the same order of magnitude.³² To illustrate how large these valuations are, we can compare our results to those from other studies. Relative to the mean valuations found in our study (of around \$200), these other studies find average valuations that are orders of magnitude smaller: \$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, 2019), and \$4.75 for information on forecasts of home prices (Fuster et al., 2018).³³

The BDM elicitation is generally preferred to the non-incentivized alternative, but it is of course not perfect – indeed, some imperfections have been documented in the literature (Shogren et al., 2001). A first concern is that our estimates of willingness to pay may be sensitive to the elicitation method – in particular, the lists of prices given in the hypothetical scenarios may act as a suggestion for what the employees “should” pay for the information. As a robustness check, we use data from a follow-up study to show that the results are robust when using an alternative (open-ended) elicitation method.³⁴ There are some additional caveats to keep in mind. One potential concern might be that answers were perceived as inconsequential, because we implemented them with only a 1% probability.³⁵ Another potential concern is that the willingness to pay could have been somewhat lower if subjects had to spend out of their pockets (Shogren et al., 2001).³⁶

Our favorite interpretation for the large willingness to pay for information about salaries

tion is 0.28 (p-value<0.001).

³²One alternative approach follows Andersen et al. (2006) by assuming that the average of the willingness to pay inside each bin is equal to the midpoint of the bin and, for the highest bin, which has no upper bound, assuming that the average is equal to the lower bound. Under this approach, the average willingness to pay is \$247 for the manager information and \$254 for the peer information. For a more parametric approach, we can use an interval regression model based on a Gaussian distribution and then use the estimates to predict the average willingness to pay. This indicates an average willingness to pay of \$333 for the manager information and \$343 for the peer information.

³³All these amounts were converted to 2017 USD PPP to be comparable to our estimates.

³⁴Results reported in Appendix C.6.

³⁵Theoretically, any positive probability of executing a choice ensures that the elicitation is incentive compatible; Carson and Groves (2007) and Charness et al. (2016) provide empirical evidence in support of this. In any case, our subjects do not have an obvious reason to distort their preferences. Other contexts, such as when asked about willingness to pay for public goods, or to evaluate a disadvantaged group, are more vulnerable to misrepresentation.

³⁶In our experiment, subjects were not paying out of pocket for the information but instead were choosing to give up money for the information. This method is used to avoid having to collect payments from subjects, and has been used often in the literature (Allcott and Kessler, 2019; Fuster et al., 2018).

is that it is not due to curiosity but primarily due to the instrumental value. For example, in the spirit of the model from Section 2.1, the employee may need the information on the manager pay to decide whether to work harder to get promoted, or to use it as bargaining chip in future salary negotiations. Indeed, Appendix C.7 provides some suggestive evidence in support of this interpretation.

3.3 Learning

If employees think they have inaccurate beliefs and they are willing to pay to acquire new information, they should incorporate that information into their belief formation once they have access to it. In this section, we measure this belief formation using a simple Bayesian learning model. We will present this model in detail, because it plays a key role in the IV estimation introduced in Section 4 below.

We follow the same econometric model that has been shown to fit the data well in information-provision experiments on a wide range of topics, such as inflation (Cavallo et al., 2017), cost of living (Bottan and Perez-Truglia, 2020), and housing prices (Fuster et al., 2018). Let subscript i index employees. Let M_i^{prior} denote the mean of the prior belief distribution about the average salary of managers – that is, the belief right before the subject reaches the information-provision experiment. Let M_i^{signal} be the value of the signal on average manager salary that we calculated for employee i (i.e., the average salary from a random sample of five managers), and let T_i^M be a binary variable that takes the value 1 if we showed that signal to employee i and 0 if not. Denote M_i^{post} as the corresponding posterior belief – that is, the perceived manager salary after the employee sees, or does not see, the information.

When priors and signals are distributed normally, Bayesian learning implies that, after the employee sees the signal, the mean of the posterior belief should be a weighted average between the signal and the mean of the prior belief, weighted by a parameter α that we refer to as the learning rate. This parameter ranges from 0 (individuals ignore the signal) to 1 (individuals fully adjust to the signal), and according to the Bayesian model it should depend on the relative precision between the prior belief and the signal. We must note this result relies on some assumptions: for example, the variance of the prior and the variance of the signal must be independent of the mean of the prior (for a formal treatment see Hoff, 2009).³⁷ We can summarize this prediction with the following equation:

$$\log(M_i^{post}) - \log(M_i^{prior}) = \alpha \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \quad (5)$$

³⁷To appreciate the significance of this assumption, consider the potential for noisy elicitation. In that case, the individuals prone to errors (noise) may also report the most extreme priors, and hence would artificially appear like they are updating by the largest extent.

In other words, the Bayesian model predicts that the belief updates should be a linear function of the gap between the signal and the prior belief. That is, respondents who over-estimated salaries would revise their beliefs downward when shown the signal, while those who underestimated salaries would revise their beliefs upward when shown the signal. In practice, there may be “spurious” reasons for individuals to have revised their beliefs in the direction of the feedback, even if the feedback had not been shown to them. Respondents may, for example, take some additional time to think when asked a question a second time and may get closer to the truth as a result. Employees may also have made a typo on their first try, which they can correct when given the chance.

The learning results for manager salary are presented in Figure 2.a. This is a binned scatterplot of the relationship between the belief revisions (y-axis) and prior gaps (x-axis). Intuitively, the x-axis shows the maximum revision we would expect if the respondent were to fully react to the information, and the y-axis shows the revision observed in practice. The red diamonds from Figure 2.a correspond to employees in the treatment group (i.e., those who were shown the information about manager salary), while the blue circles correspond to employees in the control group (i.e., those who were not shown the information about manager salary). As expected, there is a strong relationship between the belief revisions and prior gaps for employees in the treatment group: an additional percentage point (pp) in perception gap is associated with a revision that is 0.78 pp higher. In contrast, the relationship is much weaker among individuals in the control group: an additional 1 pp in the prior gap is associated with a revision that is 0.10 pp higher. This finding suggests a statistically significant (p-value <0.001) but economically small degree of spurious revision. Indeed, this result is consistent in terms of magnitude with other information-provision studies (e.g., Cavallo et al., 2017; Fuster et al., 2018).

To weed out these spurious reactions, we exploit the information provision experiment:

$$\begin{aligned} \log(M_i^{post}) - \log(M_i^{prior}) = & \tau + \alpha \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \cdot T_i^M \\ & + \beta \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) + \epsilon_i \end{aligned} \quad (6)$$

The parameter β picks up the spurious reversion towards the signal, while α picks up the true learning: that is, the degree of revisions caused by the information provision, above and beyond the spurious revisions. We do not expect subjects to fully update to the signal we provided ($\alpha = 1$) because it is based on a sample of five salaries and is thus subject to sampling variation. Nevertheless, since the precision of the signal is significantly larger than the precision of prior beliefs, we should expect α to be substantially above zero.³⁸

³⁸For example, for manager salary, the mean absolute error of our signal is 6.8% while the corresponding

Figure 2.b estimates equation (6), which yields the learning rate. The y-axis is still the revisions from prior to posterior beliefs, but the x-axis is the interaction between the perception gap and the treatment indicator. This interaction term plays the role of excluded instrument in the IV estimator presented in Section 4.1 below. Figure 2.b shows that the linear relationship predicted by the Bayesian model fits the data well. The slope of this relationship (0.69, SE 0.03) indicates that, when forming posterior beliefs about manager salary, the average employees put a weight of 69% on the signal provided by the experimenter and put the remaining 31% on their prior beliefs.³⁹

Let P_i^{prior} , P_i^{signal} , P_i^{post} and T_i^P be respectively the prior, signal, posterior, and treatment indicator for the average peer salary. We can apply the same logic used for the manager salary to the peer salary. Figures 2.c and 2.d are identical to Figures 2.a and 2.b except that they refer to peer salary instead of manager salary. The results suggest that individuals learned significantly from the peer salary information as well. The slope from Figure 2.d of 0.51 (SE 0.06) indicates that, when forming posterior beliefs about peer salary, employees put a weight of 51% on the signals of peer salary provided by the experimenter and the remaining weight of 49% on their prior beliefs about peer salary.⁴⁰ Last, Appendix C.8 shows that learning from the feedback was compartmentalized: e.g., individuals did not use the information about peer salary to update beliefs about the manager salary.

3.4 Information Diffusion

In this section, we measure whether the information given to one employee was shared with his or her peers. Measuring information diffusion can give us insights about the sources of the misperceptions documented above. Additionally, the information diffusion is also relevant for the interpretation of the results of the experiment: spillovers can create attenuation bias, because some individuals in the control group would have been exposed to the information.

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 salaries with each other, everyone in the group could form an exact guess for the average peer salary. Thus, a lack of information diffusion would help explain the large salary misperceptions we observe in the data. We can measure the degree of

figure for prior beliefs is 28%.

³⁹Appendix C.8 shows that this average learning masks some heterogeneity, with some individuals reacting fully to the signal while others react partially or do not react at all. On the other hand, the learning rates were similar across different subsets of the population such as between female and male employees.

⁴⁰While the learning rate is substantially above zero for both manager and peer salary, it is somewhat smaller for peer salary: 0.51 vs. 0.68, with a difference that is statistically significant (p-value<0.001). According to the Bayesian model, this would indicate that individuals thought that the manager signal was more precise, or felt more confident about their prior beliefs about peer salary.

information diffusion by measuring whether employees shared the salary information provided to them through the information-provision experiment. Let M_i^{true} denote the true average salary in the managerial position the respondent is guessing about. We define M_i^{abs} as the misperceptions of the posterior beliefs: $M_i^{abs} = \left| \frac{M_i^{post} - M_i^{true}}{M_i^{true}} \right|$.⁴¹ The regression of interest is as follows:

$$M_i^{abs} = \kappa_0 + \kappa_1 \cdot T_i^M + \kappa_2 \cdot I_i^M + X_i\theta + \epsilon_i \quad (7)$$

T_i^M is a binary variable indicating whether the individual received information about manager salary and thus is meant to capture the “direct” information provision. The variable I_i^M is intended to measure the “indirect” information provision, through other employees. I_i^M takes the value 1 if and only if i did not receive the information directly but is connected to a peer who received the information before the date when i responded to the survey (so that the information could have been shared with i before she started the survey).⁴² For instance, in the baseline specification, I_i^M is a binary variable that takes the value 1 if i ’s closest peer received the information.⁴³ Since the information that we provided to individuals was accurate, it should have lowered misperceptions. We thus expect $\kappa_1 < 0$. Under the hypothesis of information diffusion, we expect $\kappa_2 < 0$. All regressions include the same basic set of control variables (X_i): a linear time trend, the number of peers and the number and proportion of peers invited to the survey.⁴⁴

The regression results are presented in Table 1. Columns (1)–(4) correspond to misperceptions of manager salary.⁴⁵ In column (1), *Direct* is the binary variable indicating whether the respondent received information directly. As expected, and consistent with the findings from the previous section, the direct information provision has a strong negative effect on misperceptions, of 16 pp (p-value<0.001). For reference, the average of the dependent variable in the control group is 27.6 pp, so this effects amounts to a 58% reduction in misperceptions.

Column (2) includes an additional variable related to indirect information provision: *Clos-*

⁴¹In Appendix C.9 we show the results are robust when, instead of the absolute error in beliefs, the dependent variable is the error in beliefs. Moreover, we use histograms to provide a more flexible test of information diffusion.

⁴²In the baseline specification, we always define I_i^M to take the value 0 if i received the information directly. The rationale behind this specification is that if the individual received the information directly, the information received indirectly through peers is largely redundant. In Appendix C.9 we show the results are robust under an alternative specification.

⁴³The 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.

⁴⁴Appendix C.9 provides descriptive statistics for all of the main variables used for the analysis of information diffusion.

⁴⁵One disadvantage of using the manager salary is that different employees from the same peer group may receive information about different managerial positions, which may make the information diffusion more difficult. This concern, however, is not applicable to the case of peer salary. And, as discussed below, the results are robust for peer salary.

est Peer is a binary variable that takes the value 1 if the individual’s closest peer received the manager information before the individual responded to the survey. We define the closest peer as the peer who has exchanged the highest total of emails (sent and received) over the three months preceding 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. For instance, data on card swipes confirms that these individuals go to lunch together more often than with other peers.⁴⁶

If employees sometimes share salary information with peers, we would expect the coefficient on *Closest Peer* to be negative. Indeed, if closest peers always share the salary information with each other, the coefficient on *Closest Peer* could be as large as the coefficient on *Direct*. The results indicate, however, an absence of information diffusion. The coefficient on *Closest Peer* is close to zero (0.004), statistically insignificant, and precisely estimated. Moreover, the coefficient on *Closest Peer* (0.001) is substantially smaller than the coefficient on *Direct* (-0.160), with the difference being highly statistically significant (p-value<0.001). In other words, when we provide information about manager salary to one employee, that information affects her own subsequent perceptions but does not affect the perceptions of her closest peer. The lack of information diffusion is most likely due to a combination of factors. The firm’s pay secrecy rule may discourage employees from discussing salaries with coworkers. Some employees may not want to share salary information for strategic reasons (e.g., they may see information as a rivalrous asset). And employees may refrain from discussing salaries due to social norms around privacy (i.e., the “salary taboo”).⁴⁷

We provide a number of robustness checks. Columns (3)–(5) of Table 1 show the results where, instead of *Closest Peer*, we use alternative variables to capture indirect information provision. *No. Peers* measures the number of peers who received information (before the respondent completed the survey, as always). (*No. Peers* > 0) is a binary variable that takes the value 1 if at least one peer received the information. And *Share of Peers* measures the share of peers who received information. The results are the same: the point estimates are close to zero and are statistically insignificant and precisely estimated, indicating a lack of information diffusion. Columns (6)–(9) reproduce the same analysis as columns (1)–(5) but look at peer salary instead of manager salary. Again, we find robust evidence of an absence of information diffusion.

Last, Appendix C.9 complements the above experimental evidence with non-experimental

⁴⁶For employees working in the headquarters offices, we can use the swipe data to proxy whether a given pair of employees have lunch together: that is, whether the pair of employees swiped in and out of the building during lunch hours and within 30 seconds of each other. We find that, relative to her other peers, an employee is 53% more likely to grab lunch with her closest peer (18.4% vs. 12.0% for the other peers).

⁴⁷Indeed, in a follow-up study (Cullen and Perez-Truglia, 2018) we provide evidence in support of this channel.

tests of information diffusion. We show that, contrary to the predictions of information diffusion models, employees with higher centrality in the email network and employees who talk more often with coworkers do not have lower misperceptions.

4 Results: The Effects of Salary Beliefs on Behavior

4.1 Econometric Model

Let Y_i^{post} be a measure of employee i 's average effort (e.g., hours worked) in a period starting from the survey date and, in the baseline specification, ending 90 days later. Recall from Section 3.3 that M_i^{post} denotes the posterior belief about the average manager salary and P_i^{post} is the posterior belief about the average peer salary. The following equation establishes the relationship of interest:

$$\log(Y_i^{post}) = \eta_0 + \eta_{mgr} \cdot \log(M_i^{post}) + \eta_{peer} \cdot \log(P_i^{post}) + \nu_i \quad (8)$$

The parameter η_{mgr} captures the effects of manager salary. According to the model from Section 2.1, this parameter could be positive or negative depending on whether the career concerns or social preferences channel dominates. A higher managerial salary could incentivize the individual to work harder to get promoted ($\eta_{mgr} > 0$), but the higher manager salary could also demoralize the employee ($\eta_{mgr} < 0$). In turn, the parameter η_{peer} captures the effects of peer salary. According to the model from Section 2.1, both channels predict a negative effect ($\eta_{peer} < 0$): a higher peer salary mitigates the incentive to work hard for a promotion, and at the same time may demoralize the employee due to social preferences.

Obtaining causal estimates of η_{mgr} and η_{peer} is challenging. A simple regression of behavior on perceived salaries would be subject to the usual concerns about omitted variable biases. For instance, employees who are more optimistic about manager salary may be the same ones who have higher intrinsic motivation or higher ability, resulting in a spurious $\eta_{mgr} > 0$. Next, we introduce a simple IV estimator that exploits the exogenous variation in beliefs induced through the information-provision experiment. First, we provide the intuition behind the model with a simple example. Consider a pair of employees who have the same bias about perceived peer salary: both of them underestimate the actual manager salary by 20%. We then randomly assign information about the true manager salary to one of these two employees. We would expect that, relative to the individual who does not get the information, the individual who receives the information ends up with a perceived manager salary that is higher. For the sake of argument, let's assume that the individual who did not receive the information continues to underestimate the actual manager salary by 20%, but

that the individual who did receive the information reacts to the information and ends up underestimating the manager salary by just 10%. The information provision is thus equivalent to a +10% shock to the perceived manager salary. This allows us to check what happened to the behavior of this pair of employees in the months after they received the information. If the 10% shock to the perceived manager salary translates into higher effort, it would imply that perceived manager salary motivates employees. On the contrary, a negative effect on effort would imply that manager salary demotivates employees. Moreover, we can estimate η_{mgr} from this data. Again, for the sake of the argument, assume that the 10% shock to perceived manager salary causes a 2% increase in effort. We can calculate the implied η_{mgr} by taking the ratio between these two values: $\eta_{mgr} = 0.2 = \frac{2\%}{10\%}$.

The above analysis is based on a group of employees who underestimated their managers' salaries by 20%. In practice, only a small share of the sample will underestimate by around 20%, so there will not be enough statistical power to limit the analysis to this group alone. However, there is nothing special about the 20% underestimation in the calculations described above. We could repeat the analysis for individuals who underestimate by 50%, for individuals who overestimate by 20%, and so on and so forth. While we do not have enough power to estimate precisely within each of those groups, once we aggregate all of the groups we should have enough precision. That is what the following IV regression is designed for:

$$\log(Y_i^{post}) = \pi_0 + \eta_{mgr} \cdot \log(M_i^{post}) + \eta_{peer} \cdot \log(P_i^{post}) + \pi_1 \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \quad (9)$$

$$+ \pi_2 \cdot (\log(P_i^{signal}) - \log(P_i^{prior})) + \pi_3 \cdot \log(M_i^{prior}) + \pi_4 \cdot \log(P_i^{prior}) + X_i \pi_5 + \epsilon_i$$

$$\log(M_i^{post}) = \mu_0 + \mu_1 \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \cdot T_i^M \quad (10)$$

$$+ \mu_2 \cdot (\log(P_i^{signal}) - \log(P_i^{prior})) \cdot T_i^P + \mu_3 \cdot (\log(M_i^{signal}) - \log(M_i^{prior}))$$

$$+ \mu_4 \cdot (\log(P_i^{signal}) - \log(P_i^{prior})) + \mu_5 \cdot \log(M_i^{prior}) + \mu_6 \cdot \log(P_i^{prior}) + X_i \mu_7 + \xi_i^1$$

$$\log(P_i^{post}) = \nu_0 + \nu_1 \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \cdot T_i^M \quad (11)$$

$$+ \nu_2 \cdot (\log(P_i^{signal}) - \log(P_i^{prior})) \cdot T_i^P + \nu_3 \cdot (\log(M_i^{signal}) - \log(M_i^{prior}))$$

$$+ \nu_4 \cdot (\log(P_i^{signal}) - \log(P_i^{prior})) + \nu_5 \cdot \log(M_i^{prior}) + \nu_6 \cdot \log(P_i^{prior}) + X_i \nu_7 + \xi_i^2$$

Equations (10) and (11) correspond to the first stage of the IV regression. These regressions measure the effect of the information provision experiments on the posterior beliefs, and are based on the Bayesian learning model given by equation (6) from Section 3.3. The vector of additional control variables (X_i) is included to reduce the variance of the error term and thus improve the precision of the estimates. It contains the following variables: the employee's own salary (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. This model makes some functional form assumptions that are discussed (and relaxed) in Section 4 below.

To disentangle where the identification is coming from, it is easier to start with the reduced-form regression. In a nutshell, this exercise consists of estimating the slope between an outcome and the prior misperceptions for individuals in the control group (i.e., individuals who were randomly assigned not to receive information), and then estimate the same slope but for individuals in the treatment group (i.e., individuals who were randomly assigned to receive information). If the slopes are the same in the two groups, that would indicate that the salary information did not affect the outcome. If the slopes are statistically different between the treatment and control groups, that would indicate that the salary information affected the outcome of interest. The important part is that the differences in slopes can be attributed entirely to random assignment of the information.

While the differences in slopes can be attributed to random assignment, that does not prove that the effects of information can be attributed to the differences in posterior beliefs (i.e., learning). That is the key assumption made by the IV specification. In mathematical terms, the instrument exogeneity assumptions are $E \left[\left(\log \left(M_i^{signal} \right) - \log \left(M_i^{prior} \right) \right) \cdot T_i^M \cdot \epsilon_i \right] = 0$ and $E \left[\left(\log \left(P_i^{signal} \right) - \log \left(P_i^{prior} \right) \right) \cdot T_i^P \cdot \epsilon_i \right] = 0$. In plain English, we need to assume that the heterogeneity in the effects of information is truly due to the differences in prior misperceptions and not because of some unobserved factor that is correlated to prior misperceptions.

To address any remaining concerns about the econometric specification, 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 (i.e., in the days before the date of survey completion rather than in the days after the survey completion). We can estimate the same IV regression describe above, except that using Y_i^{prior} instead of Y_i^{post} as the dependent variable. Intuitively, the information-provision experiment should not affect behavior in the pre-treatment period because the individuals have not yet been exposed to the information. We thus expect the coefficients for η_{mgr} and η_{peer} to be close to zero and statistically insignificant in this falsification regression.

Given the timing of the survey, one potential source for bias for the effect of manager salary is that individuals updated their beliefs about peer salary upon learning the information on manager salary. This is probably not a significant source for concern in practice, however, for the following two reasons. First, we show that individuals did not use the information about peer salary to update beliefs about the manager salary. Since individuals did not update in this direction (i.e., learning about the manager salary from the information on peer salary), it is unlikely that they would update in the opposite direction (i.e., learning about the peer

salary from the information on manager salary). Second, even if individuals updated about the peer salary from the information on manager salary, that would only bias the estimated effects of the manager salary on behavior towards zero.

It is possible that different employees may react differently to the salaries of managers and peers, amounting to heterogeneity in η_{mgr} and η_{peer} . In that case, our estimates would identify the Local Average Treatment Effect (LATE) of perceptions (Imbens and Angrist, 1994) – that is, weighted averages of η_{mgr} and η_{peer} 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.

4.2 Main Results

The results from the IV regression are presented in Table 2 presents the results from the outlined above.⁴⁸ Each column of Table 2 focuses on a different form of behavior as the dependent variable. Each column reports coefficients from three different IV regressions. For now we focus on the top panel, titled *Post-Treatment (3-Months)*, in which the dependent variable is the average behavior during the 90 days after the completion of the survey.

The main outcomes of interest, effort and performance, are presented in columns (1) through (3).⁴⁹ Column (1) corresponds to the daily average number of hours worked, which is available for 29% of the sample (i.e., for employees based in the headquarters). The coefficient on $\text{Log}(\text{Manager-Salary})$ is positive (0.150) and significant, both statistically (p-value = 0.042) and economically. This coefficient indicates that believing that their managers' salaries are higher, on average, motivates employees. Since the right-hand-side and left-hand-side variables are defined in logs, this coefficient implies a behavioral elasticity of 0.150: i.e., increasing the perceived manager salary by 10% would increase the number of hours worked by 1.5%.

The effects on the other measures of effort and performance are similar to the effects on hours worked. In column (2) of Table 2, the dependent variable is the average number of emails sent, which is available for the entire subject pool.⁵⁰ The coefficient on $\text{Log}(\text{Manager-}$

⁴⁸For the sake of brevity, this table presents the IV coefficients directly. Appendix C.10 presents a break down of the results by reduced form and first stage regressions.

⁴⁹In the baseline specification, the dependent variables are in logs to accommodate proportional effects. Appendix C.11 shows that the results are roughly similar when using an alternative specification that does not require the logarithmic transformation; it also shows that the results are robust when including additional control variables: other employee characteristics, flexible controls for the prior gaps in beliefs, and higher moments of the distribution of posterior beliefs.

⁵⁰This measure of effort focuses on the total number of emails sent. In Appendix C.12, we break down the results by emails sent and received, by emails sent inside and outside of the firm, and by emails sent to

Salary) is positive (0.130) and highly significant ($p=0.001$).⁵¹ The findings are not only qualitatively consistent across the two measures of effort, but also quantitatively similar: we cannot reject the null hypothesis that the coefficient on manager salary for hours worked (0.150) is equal to the corresponding coefficient for emails sent (0.130) – $p\text{-value}=0.816$.⁵² Column (3) of Table 2 uses the measure of performance as the dependent variable, which is available for employees who have a sales role (38.4% of the sample). The coefficient on $\text{Log}(\text{Manager-Salary})$ from column (3) is positive (0.106) and on the same order of magnitude as the corresponding coefficients for the number of hours and the number of emails. Still, we should take this finding with a grain of salt because this coefficient is less precisely estimated and is thus statistically insignificant ($p\text{-value}=0.383$).

As a benchmark for the coefficients on manager salary, we turn to the coefficients on peer salary. In column (1), for the number of hours worked, we find a coefficient on $\text{Log}(\text{Peer-Salary})$ that is negative (-0.943) and statistically significant ($p\text{-value} = 0.045$).⁵³ This coefficient, equivalent to a behavioral elasticity of -0.943, is economically significant: increasing the perceived peer salary by 10% would decrease the hours worked by 9.43%. The results in columns (2) and (3) suggest that the coefficient on peer salary for hours worked (-0.943) is qualitatively and quantitatively similar to the corresponding coefficients for the number of emails (-0.431, $p\text{-value} = 0.041$) and sales performance (-0.731, $p\text{-value} = 0.014$).

One of the most important and robust findings is that the coefficients on manager and peer salary have opposite signs: while the manager salary motivates employees, the peer salary demotivates them. To provide a more rigorous comparison, the bottom rows in Table 2 report the p -value of the test of the null hypothesis that the coefficient on peer salary is equal to the coefficient on manager salary. We always reject this null hypothesis, with $p\text{-value}=0.026$ for hours worked (column (1)), $p\text{-value}=0.007$ for emails sent (column (2)), and $p\text{-value}<0.001$ for sales performance (column (1)).⁵⁴

One unique aspect of our setting is that subjects are in a continuing contract with the

employees with higher, same or lower pay grade.

⁵¹Appendix C.13 explores the heterogeneity of this effect by employee characteristics such as gender. We do not find any statistically significant evidence of heterogeneity. However, due to the precision of the coefficients, we cannot rule out considerable differences either.

⁵²This test assumes independence of the two coefficients.

⁵³The manager coefficients are substantially more precisely estimated than the corresponding peer coefficients – in column (1), for instance, the coefficient on manager salary has a standard error of 0.074 while the corresponding coefficient on peer salary has a standard error of 0.472. This difference in precision arises from the fact that our information shocks induced more variation in manager perceptions than in peer perceptions: i. prior beliefs about manager salary were less accurate than the prior beliefs of peer salary; ii. the learning rate was higher for manager salary than for peer salary.

⁵⁴The fact that the coefficients on $\text{Log}(\text{Peer-Salary})$ are larger in absolute value than the coefficients on $\text{Log}(\text{Manager-Salary})$ does not imply that horizontal comparisons are more consequential, because there is much more variation in $\text{Log}(\text{Manager-Salary})$ than in $\text{Log}(\text{Peer-Salary})$.

firm, which allows us to 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 2. Columns (4) and (5) explore two forms of retention. Column (4) uses a binary dependent variable indicating whether the employee leaves the firm in the 90 days after the survey completion. The results suggest that a 10% increase in perceived peer salary increases the probability of leaving the company by 2.32 pp (p-value = 0.029). This effect is at least directionally consistent with the effects on effort and performance: a higher perceived peer salary demoralizes employees to the extent that they are more likely to leave the firm. With regard to vertical comparisons, a 10% increase in perceived manager salary decreases the probability of leaving the company by 0.16 pp, but the effect is economically and statistically insignificant. In column (5), we use a binary dependent variable indicating whether the individual is transferred to another unit within the firm. Even though the signs of the coefficients are consistent with those in column (4), the coefficients are closer to zero and statistically insignificant. The last two columns explore effects on two forms of career progression: the logarithm of the base salary (column (6)) and an indicator variable indicating whether the individual changed positions (column (7)). In column (6), the dependent variable is the employee’s own (log) salary as of 3 months later. The coefficients on both manager and peer salary in columns (6) and (7) are close to zero, statistically insignificant, and precisely estimated, suggesting that salary perceptions did not affect these outcomes. However, these results must be taken with a grain of salt because few employees experience those career changes in such a short time horizon.

4.3 Robustness Checks

Next, we provide a number of robustness checks for the results reported above. A first concern with IV estimation is that of weak instruments (Stock et al., 2002). Given the strong reaction to the information documented in Section 3.3, this should not be a source of concern. For a rigorous assessment, Table 2 reports the Cragg-Donald F statistic, which is commonly used to diagnose weak instruments. The value of this statistic in each regression is well above the rule of thumb of $F > 10$ that was proposed by Stock et al. (2002); it takes the values of 29.8, 204.0, 98.2, 203.5, 203.4, 203.6 and 203.3 respectively in columns (1) through (7) of Table 2.

In the baseline results, we estimate the effects on behavior during the 90 days after the survey completion. We can estimate whether the effects persisted in a longer time window. It is possible that the effects weaken over time, for instance, if employees forget the information provided to them or if they revise their beliefs in light of new information. With the data we have available, we can look at the effects on effort and performance up to 180 days after the survey completion. These results are presented in the second panel of Table 2, titled *Post-*

Treatment (6-Months). A comparison between the panels *Post-Treatment (3-Months)* and *Post-Treatment (6-Months)* indicate that the coefficients are less precisely estimated in the longer time window. The point estimates indicate that, if anything, the motivating effects of manager salary gets a bit weaker over time while the demotivating effects of peer salary get a bit stronger. For a more formal test, the bottom rows of Table 2 report the difference tests between the effects computed in the 90-day and 180-day windows. We cannot reject the null hypothesis of equal coefficients across the two time windows 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 Section 2.2 we showed that, consistent with successful random assignment, the observable characteristics are balanced across the four treatment groups. The bottom panel of Table 2, titled “Pre-Treatment (Falsification)”, presents a more direct check. Those coefficients are estimated in a regression with the pre-treatment behavior (i.e., the average during the months before the survey) instead of post-treatment behavior (i.e., during the months after the survey) as the dependent variables. We expect these falsification coefficients to be close to zero and statistically insignificant: the information that was randomly provided on the date of the survey could not have possibly affected the behavior prior to the survey date. As expected, all of the falsification coefficients are close to zero and statistically insignificant. For example, the post-treatment coefficient in column (1) of Table 2 is positive (0.150, p-value = 0.042), while the corresponding pre-treatment coefficient is close to zero (0.001) and statistically insignificant.

We conducted a number of additional robustness checks that are reported in the Appendix. Two of these checks are particularly important and for that reason we summarize them below. First, in the baseline model, we make a functional form assumption that the relationship between salary perceptions and behavior is log-log linear. This is the simplest possible specification and thus provides a good starting point. This is also a common specification in the literature on relative income concerns (see e.g., Senik, 2004; Clark et al., 2008; Clark and Senik, 2010). In Appendix C.14, we use binned scatterplots to demonstrate that this log-log linear specification fits the data well. Additionally, these results show that the findings are not driven by outliers.

Another functional form assumption from the baseline model relates to the symmetry of the responses. Let us start with the main object of interest: the vertical comparisons. Our baseline specification assumes that the effects of updating beliefs upwards are the mirror image of updating beliefs downwards. However, these effects could be asymmetric in practice: e.g., finding out that the managers are paid more than initially thought may have stronger

or weaker effects than finding out that the managers are paid less than initially thought.⁵⁵ Appendix C.15 presents the results from an econometric specification that allows for this type of asymmetric responses. We do not find any significant evidence for this type of asymmetry. This evidence indicates that the symmetric specification used in the baseline model constitutes a reasonable approximation. However, due to power limitations, we cannot rule out small or moderate asymmetries either. Likewise, we do not find any evidence of this type of asymmetry (upwards vs. downward revisions) for the horizontal comparisons.

For the horizontal comparisons, there is a second type of asymmetry that may arise. Indeed, there is evidence from related studies in support of this type of asymmetry: while retention may drop when individuals are paid less than the average peer, retention does not go up as much when individuals are paid more than the average peer (Card et al., 2012; Dube et al., 2019; Breza et al., 2018). Appendix C.15 presents the results using an econometric specifications that allows for this type of asymmetric response in horizontal comparisons. When looking at the retention outcome, we find the exact same type of asymmetry reported in the related studies. When looking at the effort and performance outcomes, however, we do not find any significant evidence of asymmetric responses.⁵⁶

4.4 Mechanisms: Career Concerns

The career concerns channel suggests that when employees find out that their managers earn more, they work harder because they want to be promoted to that position. The fact that we find positive effects of manager salary on effort is directionally consistent with the career concerns channel. To probe this mechanism further, we provide two tests.

The first test is based on the prediction of the model of Section 2.1 that, to the extent that they aspire to be promoted to the managerial position, learning that the manager is better paid should make employees more optimistic about their own salary potential. To test this hypothesis, we estimate the effects of the perceived manager salary on the expectations of future salary as elicited in the survey. The results are presented in Table 3. Each of the five columns corresponds to a different dependent variable, based on the five post-treatment questions included in the survey. All coefficients are estimated with the same IV specification from Table 2.⁵⁷

⁵⁵For example, employees may have more flexibility to adjust their effort upwards in response to good news about the manager pay (e.g. by working extra hours) than to adjust their effort downwards in response to bad news.

⁵⁶The statistical power available to conduct this type of analysis is limited, so one should not conclude the effects are perfectly symmetric. However, the results do indicate that the baseline (symmetric) specification constitutes a reasonable approximation.

⁵⁷The only difference is that, when using survey outcomes, we do not observe pre-treatment outcomes so we cannot use them for falsification tests or include them as control variables. To compensate for this lack of

The dependent variable in column (1) of Table 3 is the (log) expected future salary one year in the future. Since it is highly unlikely that respondents will be promoted to the manager’s position within a year, we should expect the information on manager salary to have little or no effect on these short-term expectations. Indeed, the coefficient on *Log(Manager-Salary)* from column (1) is positive (0.025) but close to zero, statistically insignificant, and precisely estimated. Column (2) of Table 3 corresponds to the effects on the (log) expected future salary in five years instead of just one year. According to our survey data, the average employee expects a probability of 55.8% of being promoted to the manager’s position within the following five years. As a result, we would expect the information on manager salary to be more relevant in this longer time horizon. Indeed, the coefficient on *Log(Manager-Salary)* from column (5) is positive (0.166), precisely estimated, and highly statistically significant (p-value = 0.003). A 10% increase in perceived manager salary increases the expected salary in five years by 1.66%. Moreover, the magnitude of this effect (coefficient of 0.166) is not only consistent in sign but also similar in magnitude to the effects of manager salary on effort (e.g., coefficients of 0.150 for hours worked and 0.130 for emails sent).

Table 3 reports the coefficients on *Log(Peer-Salary)* too. These coefficients are also positive, but less precisely estimated and, thus, statistically insignificant: 0.071 (p-value=0.431) with respect to salary one year in the future (column (1)) and 0.280 (p-value=0.111) with respect to salary five years in the future (column (2)). This evidence would also be consistent with the career concerns channel in the model of Section 2.1, according to which employees expect to use the information on peer salary as bargaining chip.⁵⁸ Since the estimates are imprecisely estimated, however, these results must be taken with a grain of salt.

The second test of the career concerns channel is based on the prediction from the model of Section 2.1 that the effects of perceived manager salary should be stronger for managerial positions the employee could realistically aspire to attain. To conduct this test, we leverage the heterogeneity in the distance between the employee’s own position and the managerial position we asked them about. According to the survey responses, the average subject believes to need 3.65 promotions to reach the managerial position they are asked about and that there is a 55.8% probability of being promoted to that position within the next five years.⁵⁹ Most importantly, there is significant variation across subjects in how close they feel to the managerial position.

The results from this test are presented in Table 4. These specifications are identical to

pre-treatment controls, we include some additional control variables: dummies for sales role, pay band, unit, and position title.

⁵⁸This evidence is also consistent with the “tunnel effect” from Senik (2004).

⁵⁹When compared to some objective data on career progression at this firm, these perceptions seem reasonably calibrated – for details, see Appendix C.3.

the baseline specifications from Table 2 except that they allow the coefficients on manager and peer salary to be different across the following two subgroups: *Closer* and *Farther*, denoting whether the managerial positions are more or less accessible to the respondent.

We present results under two alternative specifications. In the first specification, which is presented in the top panel of Table 4, we split respondents based on whether the perceived number of promotions is below or above the median value. In the second specification, reported in the bottom panel, we split subjects by whether their perceived probability of being promoted to the managerial position is below or above the median. For example, respondents in the group *Closer* expect to be promoted to the manager position, on average, with a probability of 75.5%, while respondents in the group *Farther* expect to be promoted with a probability of 18.1%.⁶⁰

We start by discussing the results from the top panel. In columns (1) and (2) of Table 4 the dependent variables are the expected future salary in one and five years, respectively. Column (1) indicates that the perceived manager salary does not affect expectations for one year in the future, regardless of the distance to the manager. The results from column (2) indicate that the manager salary increases the employees' own salary expectations five years in the future, but only for positions that are within reach: the coefficient on $\text{Log}(\text{Manager-Salary})$ is positive (0.204) and statistically significant (p-value=0.001) for the group *Closer* but smaller (0.086) and statistically insignificant (p-value= 0.349) for the group *Farther*. We find similar heterogeneity for the effects on effort and performance. Column (3), corresponding to hours worked, shows that the coefficient on $\text{Log}(\text{Manager-Salary})$ is positive (0.212) and statistically significant (p-value=0.033) for the group *Closer* but negative (-0.074) and statistically insignificant (p-value=0.424) for the group *Farther*. Column (4), corresponding to emails sent, shows a coefficient on $\text{Log}(\text{Manager-Salary})$ that is positive (0.170) and statistically significant (p-value=0.001) for the group *Closer* but smaller (0.019) and statistically insignificant (p-value=0.856) for the group *Farther*. Column (5), corresponding to the sales outcome, the point estimate is positive (0.195) although borderline insignificant (p-value=0.137) for the group *Closer* and close to zero (0.033) and statistically insignificant (p-value=0.907) for the group *Farther*.

As a robustness check, the bottom panel of Table 4 reports the results under the alternative definition of *Closer* and *Farther*. The results from this second specification are broadly consistent with the first specification: the effects of manager salary are larger in magnitude and more statistically significant for the group *Closer* than for the group *Farther*. In summary, the evidence from Table 4 shows that, consistent with the career concerns channel,

⁶⁰The splits from Specification 1 and Specification 2 are similar but far from identical: the correlation between the expected probability of promotion and the expected number of promotions is -0.415 (p-value<0.001).

the effects of manager salary are driven by the managerial positions that the employees can aspire to attain. These differences in coefficients between *Closer* and *Farther* must be taken with a grain of salt, however. On the one hand, as reported in the bottom rows of Table 4, although large in magnitude, the differences between pairs of coefficients are not estimated with enough precision to be statistically significant. On the other hand, the fact that these differences are so robust across specifications and across outcomes suggests that they are meaningful.

4.5 Mechanisms: Social Preferences

The fact that the manager salary motivates employees goes against the predictions of the social preferences mechanism, according to which employees will be demoralized by the size of the manager’s paycheck. It is possible that social preferences are still at play but are outweighed by the motivating effects of career concerns. Below we leverage the survey outcomes to probe this mechanism more directly.

Columns (3)–(5) of Table 3 show the effects of salary perceptions on the proxies for employee morale (pay satisfaction and job satisfaction) and tolerance for inequality. Finding effects on any of these outcomes would constitute suggestive evidence in favor of the social preferences channel. Column (3) shows that the effect of $\text{Log}(\text{Manager-Salary})$ on pay satisfaction is close to zero (-0.015), statistically insignificant (p-value=0.906), and precisely estimated. Columns (4) and (5) show that the corresponding effects on job satisfaction and tolerance for inequality are also close to zero (-0.086 and 0.008), statistically insignificant (p-value=0.399 and 0.920), and precisely estimated. These three coefficients imply that a 10% increase in manager salary would reduce pay satisfaction by a mere 0.16% of a standard deviation, job satisfaction by 1.1% of a standard deviation, and tolerance for inequality by just 0.14% of a standard deviation.

In contrast to the findings for manager salary, we find evidence that the peer salary has significant effects on these same three survey outcomes. In column (3), the effect of peer salary on pay satisfaction is negative (-0.762) and statistically significant (p-value=0.078). This effect is economically large, implying that a 10% higher peer salary decreases pay satisfaction by roughly 8.3% of a standard deviation. Moreover, we can reject the null hypothesis that the coefficient on peer salary (-0.762) is equal to the coefficient on manager salary (-0.015), with a p-value=0.084. The effect of peer salary on job satisfaction is statistically insignificant (p-value=0.366), but still negative (-0.444, from column (4)) and similar in magnitude and statistically indistinguishable from the corresponding effect on pay satisfaction (-0.762, from column (3)). The peer salary coefficient from column (5) is negative (-0.373) and statistically significant (p-value=0.084). This coefficient implies that a 10% increase in peer salary reduces

tolerance to pay inequality by 6.5% of a standard deviation.

The above findings suggest that the social preferences channel is at play for horizontal comparisons but not for vertical comparisons.⁶¹ This result is stunning in light of how much larger the vertical salary gaps are, relative to gaps in peer salary. We think there are three alternative interpretations for this finding. The first interpretation relates to reference groups. A large body of research on social comparisons argues that individuals tend to compare themselves not to everyone else, but to other people in specific reference groups (Clark and Senik, 2010). In our context, employees may find natural to compare their own salaries to the salaries of other employees in the same position but may not feel inclined to compare themselves to their managers. A second interpretation is based on fairness concerns. More precisely, employees may feel demoralized about horizontal comparisons because they feel like everyone working in the same position should be paid the same regardless of their skill and productivity. By contrast, employees and managers have different positions and responsibilities, and as a result the vertical differences are not seen as unfair.

There is a third interpretation, also based on fairness concerns. Employees may be intolerant to the horizontal salary differences because they perceive those differences arise largely due to non-meritocratic factors such as luck and office politics. Indeed, the findings from Breza et al. (2018) suggest that small horizontal differences in pay can be justified if they arise due to large differences in productivity. By contrast, employees may find it a lot easier to justify the vertical salary differences as meritocratic. For example, employees may think that managers deserve higher salaries because they add more value to the firm or because they worked hard to get to that position. Indeed, this interpretation echoes a robust finding in the literature about preferences for redistribution that some poor people do not want to tax the rich because they think the rich are deserving of their wealth (Di Tella et al., 2016).

While we do not have direct evidence on whether employees perceive horizontal salary differences to be meritocratic or not, we do have some suggestive evidence that non-meritocratic factors may play a significant role in this context. Several factors can be at play when determining horizontal salary differences. While there are rules governing the minimum and maximum salaries for a given position, salaries are determined at the individual level and there is a lot of room for negotiation and discretion.⁶² In principle, some of these factors may be perceived as meritocratic, while other factors may be perceived as non-meritocratic. For example, some employees may perceive that some of their peers are paid more because

⁶¹While the above evidence suggests that social preferences may play a significant role in horizontal comparisons, we by no means suggest that this is the only or even the main mechanism at play. Since studying horizontal comparisons is not the main contribution of this paper, we defer the discussion of other potential mechanisms to Appendix C.16.

⁶²For a detailed description of how salaries are determined at this organization, see Appendix C.17.

they were high performers in recent years. Other employees, however, may perceive that the highest-paid peers are being unfairly favored by their manager. For a more quantitative assessment, Appendix C.17 shows that while employees who work harder tend to earn more than their peers, the strength of that relationship is quite small in magnitude. Although speculative, we take this as evidence that employees may perceive horizontal differences to be largely non-meritocratic.

5 Discussion

5.1 Generalizability of the Results

In this section, we discuss how the results from this setting can be generalized to other settings.

Following List (2020), we start with the *litmus test* for external validity by describing the ideal setting. To study how employees learn about the salaries of their managers and how they react to that information, the setting should satisfy three important criteria. First, there must be a clear hierarchy in the organization where employees are regularly promoted to higher positions. Second, employees should be employed in long-term contracts with well-defined base salary. Third, pay-setting practices, such as transparency, performance pay and inequality, should be representative of most large firms. All these three features are present, arguably, in most white-collar jobs around the world. Our field experiment could have been conducted not only in this Southeast Asian bank, but in any of those other firms, from a tech company in Silicon Valley, to a law firm in Jakarta and a hospital from Brazil. In theory, the economic incentives should be the same regardless of whether the firm is in the banking or health care sectors, or whether the employees are selling credit cards or performing surgeries.

So, why did we conduct this experiment in this specific firm? This is where we enter the realm of feasibility. Our empirical test demanded that the firm was large enough to be able to recruit thousands of employees for the experiment. The test demanded that it was possible to use administrative records to track the behavior of the employees, such as their effort and performance, and to observe sensitive information such as the employees' salaries, the organization structure, and the network between co-workers. Last and most importantly, the empirical test requires giving permission to researchers to ask sensitive questions and reveal sensitive salary information in a controlled manner. These are the reasons that made our firm the ideal setting for the empirical test.

As discussed in Section 2.2, the firm in which we conducted the study is comparable to other large firms around the world in some key respects such as pay inequality and pay

transparency. This does not mean that we expect the results from this study would be identical if conducted in a different firm. On the contrary, according to the economic theories we test for, the results should be different depending on some mediating factors.

First, consider the case of career concerns. According to that mechanism, learning that the managers are well-paid should be motivating only to the extent that the employee can expect to one day be in the shoes of the manager. Indeed, we provide direct evidence that, even within the employees of this firm, the motivating effects were significant only for the employees who could aspire to attain the managerial position. In other firms, the opportunities to climb up the ladder may be negligible, and so should be the motivating effects of vertical comparisons. For example, the drivers in ride-sharing apps cannot aspire to be promoted to management positions; and floor workers in some garment factories may be expected to forever remain at the bottom of the ladder. As a result, we would not expect those workers to react positively to the news that their managers' pay is higher than they thought. Next, consider the case of social preferences. According to that mechanism, learning that one's managers are well-paid may be demotivating if employees think that is unfair. We study a private sector organization operating in a competitive financial industry and the general view inside the organization seems to be that promotions are largely performance based. If, instead, we were studying public employees in a corrupt government, or offspring operating a family business, then the disclosure of large differences in pay between employees and management could have elicited very different reactions, including anger and resentment.

There are some differences across countries, such as in the levels of inequality, culture or social norms, that may mediate the effects that we study and should then be taken into consideration while extrapolating our findings to other settings. Regarding the career concerns mechanism, we do not see any obvious reasons to expect culture to matter. For the social preferences channel, however, there are reasons to believe cultural differences could matter. For example, some differences in the strength of social preferences have been documented across countries, although those cross-country differences tend to be much smaller than the corresponding differences within individuals of the same country (Falk et al., 2018).⁶³ Since we are the first to study vertical comparisons, future research is needed to shed light on the role of culture and social norms as mediating factors. Regarding horizontal comparisons, however, there is already evidence from a variety of settings. For example, Card et al. (2012) conducted their experiment with employees from a large university in the United States; Breza et al. (2018) conducted their experiment with manufacturing workers from a small Indian firm; Cohn et al. (2014) conducted their experiment with sales representatives from a

⁶³The similarity of social preferences across countries is also consistent with the view held by anthropologists that social preferences are human universals (Brown, 1991).

small German firm; and [Dube et al. \(2019\)](#) leverage a quasi-experiment involving entry-level workers from a large retail firm in the United States; and our own experiment was conducted with employees from a large financial institution from Southeast Asia. Despite the settings involve different firms in different countries and industries, the finding that horizontal inequality has negative effects on employee morale is quite robust across all of those different settings.

To aid the reader in extrapolating our findings to other settings, we provide a number of descriptive statistics about the firm and the subjects. For instance, we document that a majority (73%) of the employees in our subject pool are female. If, say, information frictions were more important for female employees than for male employees, then our findings would probably over-estimate the extent of information frictions in settings with a minority of female employees. Indeed, we break down our findings by gender. To the extent that we do not find any statistically significant evidence by gender, the gender composition may not be too important. However, some of the heterogeneity has limited statistical power so those results must be taken with a grain of salt.⁶⁴

We can also evaluate the generalizability of the results of our field experiment by following the SANS conditions (Selection, Attrition, Naturalness, and Scalability) from [List \(2020\)](#). In terms of selection, the subject pool is highly representative of the target population.⁶⁵ This is an experimental context in which attrition is not an issue because the outcomes can be tracked with administrative records. In terms of naturalness of the setting, we conduct a natural field experiment that is ideal in this dimension ([Harrison and List, 2004](#)). The natural, high-stakes setting is particularly important for the type of economic theories being tested, social preferences and career concerns [Levitt and List \(2007\)](#); [Al-Ubaydli and List \(2013\)](#). Last, this is a context for which there are no concerns regarding scalability. Given the nature of our intervention, providing simple information to the subjects, it should be straightforward to scale the intervention to the universe of employees. And due to the nature of the intervention too, there are no obvious reasons to expect a “voltage drop” after scaling up the intervention ([Al-Ubaydli et al., 2017](#)).

Using the terminology of [List \(2020\)](#), we intend this paper to be a Wave 1 study, which focuses on establishing initial causality and producing first tests of theory. Although we offer evidence from a specific firm, our research design can be adapted to other firms. As discussed above, we expect that the results may change depending on some mediating factors such as the organizational structure and cultural factors. Thus, we hope our research design will be

⁶⁴Appendix [C.5](#) presents the heterogeneity in misperceptions and willingness to pay, Appendix [C.8](#) presents the heterogeneity in learning, and Appendix [C.13](#) presents some heterogeneity in the effects of perceptions on behavior.

⁶⁵For details, see Appendix [C.3](#).

applied in other firms and in other countries, to shed light on the mediating factors and to provide further support of the underlying economic models.

5.2 Implications for Pay Transparency

Our findings relate to the choice of pay transparency policies by firms as well as the regulation of those policies by the government.

The revealed-preference evidence from Section 3.1, that some employees are willing to pay significant amounts to acquire salary information, suggests that some employees have much to gain from having access to more information about salaries. To provide complementary evidence on this matter, we included two questions on preferences for transparency in our employee survey. We explained that salaries are currently confidential information at the firm, and asked employees how they felt about alternative disclosure policies. In the first scenario, we propose the creation of a website showing the same type of information that we provided in our field experiment: i.e., the average salaries by position and unit.⁶⁶ The responses indicate that a majority of employees (65.26%) favor the policy, while 14.22% feel indifferent and only a minority (20.52%) opposes it. This finding that employees report a desire for more pay transparency has been documented in other firms too – e.g., a survey of employees from eight developed countries shows that most employees wish their employers were more transparent about pay (Glassdoor, 2016). It turns out, however, that whether employees support higher transparency or not depends crucially on how the information is disclosed. In a second scenario, we offer employees to replace the status quo policy by a website that shows itemized information about salaries. In other words, you can use this website to look up any specific employee and find out how much they get paid, and other employees can look up your salary too. There is little support among employees for the disclosure of itemized information: a strong majority (74.83%) opposes this disclosure policy, while 11.84% feel indifferent and only a minority (13.33%) supports it. One plausible interpretation is that while employees value the salary information a lot, they may value their privacy even more (Cullen and Perez-Truglia, 2018).

From the firm’s perspective, however, it is less obvious whether pay transparency would improve or worsen the outcomes that the firm cares about, such as effort and performance. Our information provision experiment was not designed with the goal of increasing average effort and performance, but to induce shocks to perceived salaries. However, the findings from Section 4 do suggest that firms could gain from being more transparent about pay. For horizontal transparency, we documented that the perceptions about peer salary are, on

⁶⁶Employees could report their support or opposition to this new policy using the following scale: “strongly in favor,” “in favor,” “I would not care,” “against,” or “strongly against”. For more details, see Appendix C.18.

average, accurate. As a result, disclosing information about peer salary does not change the average perceived peer salary and, thus, would not affect the average effort either. For vertical transparency, however, the perceived manager salary is, on average, under-estimated. As a result, disclosing information on manager salary may increase the average perception of manager salary and, in turn, result in small gains in average effort. In Appendix C.19 we provide direct estimates of the average effects of disclosing information, which are directionally consistent with this prediction reported above, although imprecisely estimated. Moreover, the results from Section 4 provide hints for ways in which firms can design their transparency policies to boost employee motivation. For example, our evidence suggests that the effects of vertical comparisons are strongest when employees learn about managerial positions they aspire to attain. In our experiment, we disclosed salary information for a single managerial position that may or may not be the most relevant for the employee. Instead, firms may want to ask their employees to list all the managerial positions they aspire to attain, and then provide feedback on the average salaries in those specific positions. This hypothetical intervention would probably have a stronger impact on the average effort than the non-targeted information disclosed in our experiment.

Our results relate to some of the debates around pay transparency too. The evidence of significant information frictions presented in Section 3.1 indicates that policies that mandate pay disclosure may go a long way towards making sure that employees have access to salary information. There is, however, a second widespread motivation for pay transparency mandates: they may put pressure on firms to lower pay inequality. For example, some firms even experimented with capping the ratio of earnings at the top and bottom with the explicit intention of increasing employee morale (Dvorak, 2007). Our evidence suggests that this argument is valid, but in a narrow sense: while transparency may pressure firms to reduce horizontal inequality, our findings suggest that employees are unlikely to put pressure to reduce vertical inequality.⁶⁷ To the extent that vertical inequality makes up the vast majority of the pay inequality at firms, our evidence suggests that transparency policies may not be so effective at curbing inequality.

However, as discussed in Section 5.1 above, it is important to keep in mind that our results are based on a single firm from a specific country. Thus, we must be cautious when extrapolating the results from our setting to other settings. Moreover, there are some specific features of our transparency intervention that may be important to keep in mind too. For

⁶⁷It is still possible that these policies are effective by exposing the inequality in the general public. However, if employees inside the firm find the vertical inequality acceptable, individuals outside of the firm may find it acceptable too. Also, our results indicate that pay transparency may still be useful to induce horizontal pay compression, such as paying employees the same within a given position, or paying men and women equally for the same job.

instance, our experiment did not provide employees with information about managerial pay at the very top of the organization (e.g., informing tellers about the salary of the CEO). When learning about the salaries of top executives, the vertical inequality may be so large that some employees may find it unfair – however, our data cannot speak to that conjecture. This is an important caveat to keep in mind because some pay transparency policies are aimed precisely at disclosing the salaries of the top executives.

6 Conclusions

We presented the results from a field experiment involving 2,060 employees from a multibillion-dollar corporation. The research design combines survey data, administrative data, and an information-provision experiment to shed light on how employees learn about the salaries of their managers and peers, and how those beliefs affect their own behavior. We documented large misperceptions of the salaries of managers and peers. We showed that perceptions of the salaries of managers and peers have significant effects on the employee’s own behavior. When they find out that their managers earn more than they thought, employees tend to work harder. On the contrary, employees are demotivated when they find out that their peers earn more. Additionally, we provided suggestive evidence of two causal mechanisms at play: career concerns and social preferences.

Our findings have a number of implications. For instance, our findings suggest that changing the salary of one employee can affect the behavior of other employees in the same firm. Thus, these externalities should be taken into account when designing compensation incentives inside organizations. Indeed, our findings provide a new perspective on how firms set compensation. 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 the manager level has a positive externality on the behavior of all employees who aspire to be promoted to that level. Because of these externalities, firms may find it optimal to load rewards vertically rather than horizontally. Indeed, these findings may help to explain why firms tend to provide financial incentives vertically, in the form of promotions, rather than providing horizontal incentives such as pay-for-performance (Baker et al., 1988).

Our findings have implications for pay transparency too. Our results suggest that most employees would be better off if their employers were more transparent about pay and, moreover, transparency can help motivate employees. Our evidence is also related to a recent debate on pay transparency laws. There is a widespread view that forcing firms to be more transparent would reduce pay inequality.⁶⁸ Our findings suggest that these policies may be

⁶⁸For example, policy makers promote transparency policies, such as disclosure of CEO pay (Faleye et al.,

effective, but in a narrow sense: while transparency may pressure firms to reduce horizontal inequality, our findings suggest that employees are unlikely to exert the same pressure to reduce vertical inequality, which constitutes the bulk of pay inequality.

Last, our study discusses some mediating factors that need to be considered when extrapolating the results from this specific setting to other settings. Furthermore, we discuss avenues in which future research may extend our analysis to provide a deeper understanding of the economic forces at play.

References

- Akerlof, G. A. (1982). Labor Contracts as Partial Gift Exchange. *The Quarterly Journal of Economics* 97(4), 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.
- Al-Ubaydli, O. and J. A. List (2013). On the Generalizability of Experimental Results in Economics. In Frechette, G. & Schotter, A., *Methods of Modern Experimental Economics*, Oxford University Press.
- Al-Ubaydli, O., J. A. List, D. LoRe, and D. Suskind (2017). Scaling for Economists: Lessons from the Non-Adherence Problem in the Medical Literature. *Journal of Economic Perspectives* 31(4), 125–144.
- Allcott, H. and J. B. Kessler (2019). The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons. *American Economic Journal: Applied Economics* 11(1), 236–276.
- 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.
- Baker, G., M. Gibbs, and B. Holmstrom (1994). The Internal Economics of the Firm: Evidence from Personnel Data. *The Quarterly Journal of Economics* 109(4), 881–919.
- Baker, G., M. Jensen, and K. Murphy (1988). Compensation and Incentives: Practice vs. Theory. *The Journal of Finance* 43(3), 593–616.
- Becker, G. M., M. H. DeGroot, and J. Marschak (1964). Measuring Utility by a Single-Response Sequential Method. *Behavioral Sciences* 9(3), 226–232.
- Bottan, N. L. and R. Perez-Truglia (2020). Choosing Your Pond: Location Choices and Relative Income. *The Review of Economics and Statistics*, forthcoming.
- Bracha, A., U. Gneezy, and G. Loewenstein (2015). Relative pay and labor supply. *Journal of Labor Economics* 33(2), 297–315.

2013; Mas, 2016; Mueller et al., 2017), with the hope of putting pressure on firms to reduce pay inequality.

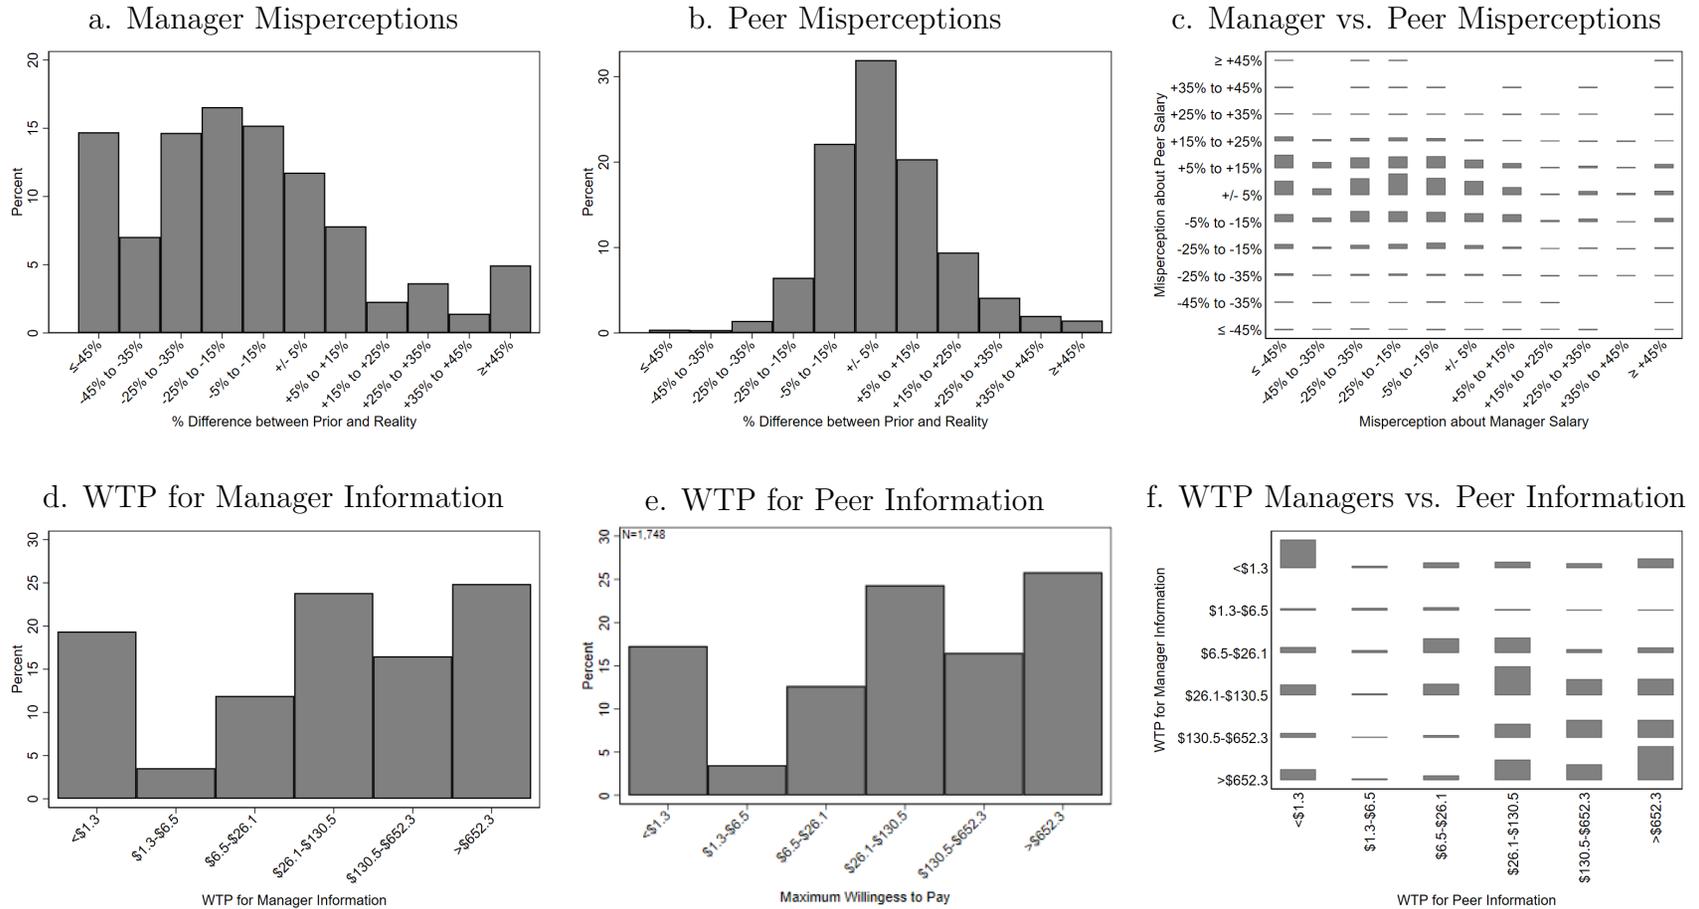
- Breza, E., S. Kaur, and Y. Shamdasani (2018). The Morale Effects of Pay Inequality. *The Quarterly Journal of Economics* 133(2), 611–663.
- Brown, D. (1991). *Human Universals Hardcover*. Philadelphia, PA: Temple University Press.
- Burchell, B. and D. Yagil (1997). Socioeconomic and Political Initiators of Pay Comparison. *Work, Employment, & Society* 11(4), 737–748.
- Caldwell, S. and N. Harmon (2018). Outside Options, Bargaining and Wages: Evidence from Coworker Networks. *Working Paper*.
- 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.
- Carson, R. T. and T. Groves (2007). Incentive and informational properties of preference questions. *Environmental and Resource Economics* 37(1), 181–210.
- 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., U. Gneezy, and B. Halladay (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization* 131, 141–150.
- 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. 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.
- CNBC (2017). A brief history of the 8-hour workday. *CNBC, May 3rd 2017*.
- 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, 877–898.
- 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). The Salary Taboo: Privacy Norms and the Diffusion of Information. *NBER Working Paper No. 25145*.
- Cullen, Z. and R. Perez-Truglia (2019). The Old Boys’ Club: Schmoozing and the Gender Gap. *NBER Working Paper No. 26530*.

- Danz, D., L. Vesterlund, and A. J. Wilson (2020). Belief Elicitation: Limiting Truth Telling with Information on Incentives. *NBER Working Paper No. 27327*.
- Day, N. E. (2007). An investigation into pay communication: Is ignorance bliss? *Personnel Review* 36, 739–762.
- Dellavigna, S., J. A. List, U. Malmendier, and G. Rao (2019). Estimating Social Preferences and Gift Exchange with a Piece-Rate Design. *NBER Working Paper No. 22043*.
- Dewatripont, M., I. Jewitt, and J. Tirole (1999). The Economics of Career Concerns, Part I. *The Review of Economic Studies* 66(1), 183–198.
- Di Tella, R., J. Dubra, and A. L. Lagomarsino (2016). Meet the Oligarchs: Business Legitimacy, State Capacity and Taxation. *NBER Working Paper No. 22934*.
- Dube, A., L. Giuliano, and J. Leonard (2019). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review* 109(2), 620–663.
- Dvorak, P. (2007). Limits on Executive Pay: Easy to Set, Hard to Keep. *The Wall Street Journal*. April 9, 2007.
- Edwards, M. A. (2005). The Law and Social Norms of Pay Secrecy. *Berkeley Journal of Employment and Labor Law* 26, 41–63.
- Faleye, O., E. Reis, and A. Venkateswaran (2013). The determinants and effects of CEO-employee pay ratios. *Journal of Banking & Finance* 37(8), 3258–3272.
- Falk, A., A. Becker, T. Dohmen, D. B. Huffman, and U. Sunde (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics* 133(4), 1645–1692.
- Frank, R. (1984). Are Workers Paid Their Marginal Products? *American Economic Review* 74(4), 549–571.
- Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *The Review of Economics and Statistics*, forthcoming.
- Gee, K. (2017). Ask Me How Much Money I Make. *The Wall Street Journal*, October 26 2017.
- 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.
- Gibbons, R. and M. Waldman (1999a). A Theory of Wage and Promotion Dynamics Inside Firms. *Quarterly Journal of Economics* 114(4), 1321–1358.
- Gibbons, R. and M. Waldman (1999b). Careers in Organizations: Theory and Evidence. In *Handbook of Labor Economics, Volume 3B*, pp. 2373–2437.
- Glassdoor (2016). Global Salary Transparency Survey. Technical report.
- Gneezy, U. and J. A. List (2006). Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments. *Econometrica* 74(5), 1365–1384.

- 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.
- Harris, M. and B. Holmstrom (1982). A Theory of Wage Dynamics. *The Review of Economic Studies* 49(3), 315–333.
- Harrison, G. W. and J. A. List (2004). Field Experiments. *Journal of Economic Literature* 42(4), 1009–1055.
- Hegewisch, A., C. Williams, and R. Drago (2011). Pay Secrecy and Wage Discrimination. *Institute for Womens Policy Research, June 9 2011*.
- 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? *Working Paper*.
- Imbens, G. W. and J. D. Angrist (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica* 62(2), 467–475.
- Jovanovic, B. (1979). Job Matching and the Theory of Turnover. *Journal of Political Economy* 87(5), 972–990.
- 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.
- Kiatpongsan, S. and M. I. Norton (2014). How Much (More) Should CEOs Make? A Universal Desire for More Equal Pay. *Perspectives on Psychological Science* 9(6), 587–593.
- Lawler, E. E. (1965). Managers' Perceptions of their Subordinates' Pay and of Their Superiors' Pay. *Personnel Psychology* 18(4), 413–422.
- Lawler, E. J. (2003). Pay Practices in Fortune-1,000 Corporations. *WorldatWork Journal* 12(4), 45–54.
- Lazear, E. P. (1989). Pay Equality and Industrial Politics. *Journal of Political Economy* 97(3), 561–580.
- Lazear, E. P. and S. Rosen (1981). Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy* 89(5), 841–864.
- Levitt, S. D. and J. A. List (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives* 21(2), 153–174.
- List, J. A. (2020). Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial. *NBER Working Paper No. 27535*.
- Mas, A. (2016). Does Disclosure affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act. *Working Paper*.

- Mas, A. (2017). Does Transparency Lead to Pay Compression? *Journal of Political Economy* 125(5), 1683–1721.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics* 99(2), 210–221.
- Mueller, H. M., P. P. Ouimet, and E. Simintzi (2017). Within-Firm Pay Inequality. *The Review of Financial Studies* 30(10), 3605–3635.
- PayScale (2018). 2018 Compensation Best Practices Report. Technical report.
- Perez-Truglia, R. (2020). The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment. *American Economic Review* 110, 1019–54.
- Romer, D. (1984). The Theory of Social Custom: A Modification and Some Extensions. *The Quarterly Journal of Economics* 99(4), 717–727.
- Rosen, S. (1986). Prizes and Incentives in Elimination Tournaments. *The American Economic Review* 76(4), 701–715.
- Scott, D. (2003). Survey of Compensation Policies and Practices.
- Senik, C. (2004). When information dominates comparison: Learning from Russian subjective panel data. *Journal of Public Economics* 88(9), 2099–2123.
- Shogren, J. F., S. Cho, C. Koo, J. List, C. Park, P. Polo, and R. Wilhelmi (2001). Auction mechanisms and the measurement of WTP and WTA. *Resource and Energy Economics* 23, 97–109.
- Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming up inequality. *Quarterly Journal of Economics* 134(1), 1–50.
- Stock, J. H., J. H. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics* 20(4), 518–529.
- Trachtman, R. (1999). The Money Taboo: Its Effects in Everyday Life and in the Practice of Psychotherapy. *Clinical Social Work Journal* 27(3), 275–288.
- Wen, Q., P. A. Gloor, A. F. Colladon, P. Tickoo, and T. Joshi (2020). Finding top performers through email patterns analysis. *Journal of Information Science* 46(4), 508–527.

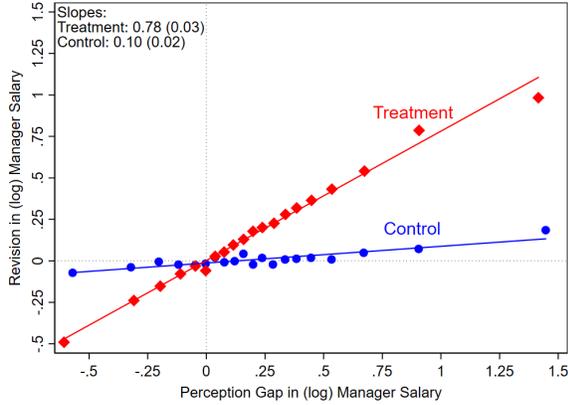
Figure 1: Salary Misperceptions and Willingness to Pay for Salary Information



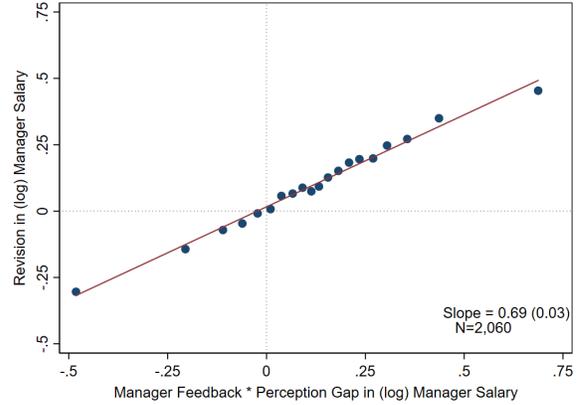
Notes: Panel (a)–(c) correspond to salary misperceptions defined as the employee’s prior belief (according to an incentivized survey question) and the actual salary (according to the firm’s administrative records), divided by the actual salary ($N=2,060$). Panel (a) is about the average manager salary, panel (b) is about the average peer salary, and panel (c) is about the joint distribution of (a) and (b). Panels (d)–(e) shows the willingness to pay (WTP) for a specific information piece (an average based on a sample of five manager/peer salary) based on the responses to multiple price list questions. Panel (a) corresponds to the WTP for the manager information ($N=1,637$ respondents with consistent responses across the five scenarios). Panel (b) corresponds to the WTP for peer information ($N=1,748$ respondents with consistent responses across the five scenarios). Panel (c) corresponds to the joint distribution of (a) and (b) ($N=1,478$ respondents with consistent responses for both the manager and peer scenarios).

Figure 2: Effects of Information Provision Experiment on Salary Perceptions

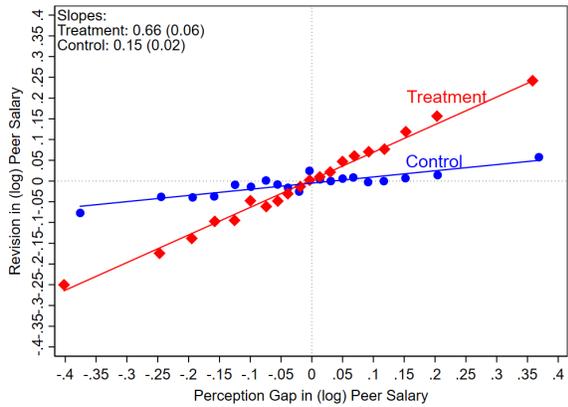
a. Manager Salary: Treatment vs. Control



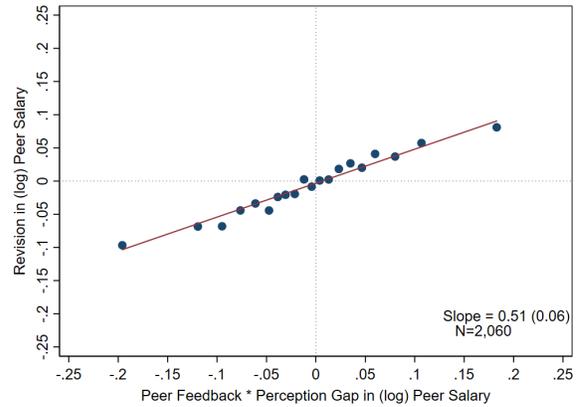
b. Manager Salary: Bayesian Learning



c. Peer Salary: Treatment vs. Control



d. Peer Salary: Bayesian Learning



Notes: $N=2,060$. In panel (a), the y-axis is the respondent’s update (i.e., the posterior belief about the average manager salary minus the corresponding prior belief) and the x-axis corresponds to the difference between the feedback chosen for the employee (the average salary among the random sample of 5 managers) and the respondent’s prior belief. The red diamonds (labeled “Treatment”) correspond to the respondents who were shown the feedback about manager salary and the blue circles correspond to the respondents who were not shown such feedback. Panel (b) estimates the Bayesian learning equation (6) from Section 3.3. The y-axis is the same as that from panel (a), while the x-axis corresponds to the same x-axis from panel (a) but multiplied by a binary 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 dots correspond to the binned scatterplot, the slope to the linear regression, and the standard error are clustered at the position level and presented in parentheses. Panel (c) and (d) are equivalent to panel (a) and (b) except that they are about peer salary instead of manager salary.

Table 1: Information Diffusion

	Misperceptions on Manager Salary					Misperceptions on Peer Salary				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Information Assignment										
Direct	-0.160*** (0.011)	-0.159*** (0.011)	-0.168*** (0.013)	-0.168*** (0.017)	-0.156*** (0.014)	-0.045*** (0.005)	-0.044*** (0.005)	-0.042*** (0.007)	-0.043*** (0.009)	-0.044*** (0.007)
Indirect: Closest Peer		0.001 (0.019)					0.009 (0.011)			
Indirect: No. Peers			-0.003 (0.003)					0.001 (0.001)		
Indirect: (No. Peers > 0)				-0.011 (0.019)					0.002 (0.009)	
Indirect: Share of Peers					0.031 (0.069)					0.012 (0.028)
Constant	0.276*** (0.019)	0.276*** (0.019)	0.276*** (0.019)	0.280*** (0.021)	0.274*** (0.019)	0.111*** (0.011)	0.111*** (0.011)	0.111*** (0.011)	0.110*** (0.012)	0.110*** (0.011)
Observations	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060

Notes: N= 2,060. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the peer group level. Each column corresponds to a different regression. All regressions follow the econometric model described in Section 3.4. In columns (1)-(5), the dependent variable is the absolute value of the difference between the posterior belief about average manager salary and the true average, divided by the true average. *Direct* is a binary variable indicating if the subject received the signal on manager salary. The other independent variables measure if the peers of the subject received the information before the subject completed the survey. Those variables take the value 0 if the employee received the information directly. *Closest Peer* is a binary variable indicating if the individual's closest peer (defined as the peer with whom the employee exchanges the most number of emails in Jan-Mar 2017) received the information. *No. Peers* is the number of peers who received the information. *(No. Peers > 0)* is a binary variable indicating if at least one peer received the information. *Share of Peers* is the share of peers who received information. All regressions include as control variables the date when the survey was completed, peer group size and the number and proportion of the peer group invited to the survey. Columns (6)-(10) are equivalent to columns (1)-(5), but using peer salary instead of manager salary.

Table 2: Effects of Salary Perceptions on Behavior

	Effort and Performance			Career Moves			
	$\log(Hours)$ (1)	$\log(Emails)$ (2)	$\log(Sales)$ (3)	P(Left) (4)	P(Transfer) (5)	$\log(Salary)$ (6)	$P(\Delta Title)$ (7)
Post-Treatment (3-Months):							
$Log(\text{Manager-Salary})^{(i)}$	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)	-0.016 (0.023)	-0.003 (0.030)	0.002 (0.011)	0.012 (0.029)
$Log(\text{Peer-Salary})^{(ii)}$	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	0.232** (0.106)	0.093 (0.106)	0.004 (0.052)	0.114 (0.123)
Post-Treatment (6-Months):							
$Log(\text{Manager-Salary})^{(iii)}$	0.112 (0.091)	0.073 (0.056)	0.233*** (0.087)	0.004 (0.030)	0.012 (0.030)	0.029 (0.048)	0.010 (0.029)
$Log(\text{Peer-Salary})^{(iv)}$	-1.200** (0.591)	-0.437** (0.209)	-1.064*** (0.347)	0.213* (0.127)	-0.007 (0.157)	-0.212 (0.176)	0.120 (0.123)
Pre-Treatment (Falsification):							
$Log(\text{Manager-Salary})$	0.001 (0.114)	-0.101 (0.071)	0.063 (0.160)		0.029 (0.029)	0.002** (0.001)	0.009 (0.010)
$Log(\text{Peer-Salary})$	-0.205 (0.542)	-0.184 (0.289)	-0.191 (0.412)		0.212 (0.163)	-0.001 (0.005)	-0.071* (0.040)
P-value $H_0: (i)=(ii)$	0.026	0.007	0.000	0.015	0.398	0.963	0.424
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
Cragg-Donald F-Stat.	29.8	204.0	98.2	203.5	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 three sets of IV regressions, following the specification from Section 4.1: in *Post-Treatment (3-Months)*, the dependent variable is the average behavior during the 90 days after the completion of the survey; in *Post-Treatment (6-Months)* the dependent variable is the average behavior during the 180 days after the completion of the survey; in *Pre-Treatment (Falsification)*, the dependent variable is the average behavior during 30 days before the completion of the survey. *Manager-Salary* is the posterior belief about the average manager salary and *Peer-Salary* is the posterior belief about the average peer salary. All regressions control for monthly lags of the dependent variable, (log) own salary, (log) tenure, 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. $P(Left)$, $P(Transfer)$ 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 (beginning) of the post-treatment (pre-treatment) period. The *Cragg-Donald F-Stat.*, *Mean Outcome* and *Std. Dev. Outcome* correspond to the regressions in the first panel (*Post-Treatment (3-Months)*). In columns (1), (2), (3) and (6), *Mean Outcome* and *Std. Dev. Outcome* correspond to the values prior to taking the logarithm function. In column (4), the pre-treatment coefficient cannot be estimated because by construction there is no variation in the pre-treatment outcome. Column (1) corresponds to the subsample of employees in the headquarter offices and column (3) to the subsample of employees with sales roles.

Table 3: Effects of Salary Perceptions on Survey Outcomes

	Log(E[Future Salary])		Rank(Prod.)	Satisfaction		Ineq. Tol.
	+1 year (1)	+5 years (2)	(3)	w/Pay (4)	w/Job (5)	(6)
<i>Log (Manager-Salary)⁽ⁱ⁾</i>	0.025 (0.025)	0.166*** (0.055)	0.000 (0.015)	-0.015 (0.125)	-0.086 (0.102)	-0.008 (0.075)
<i>Log (Peer-Salary)⁽ⁱⁱ⁾</i>	0.071 (0.090)	0.280 (0.176)	0.044 (0.040)	-0.762* (0.433)	-0.444 (0.491)	-0.373* (0.216)
P-Value (i)=(ii)	0.595	0.532	0.280	0.084	0.433	0.135
Cragg-Donald F-Stat.	253.5	255.3	250.5	253.6	254.3	254.3
Mean Dep. Var.	2.58	3.22	0.47	2.79	3.60	1.80
Std. Dev. Dep. Var.	0.51	0.59	0.22	0.92	0.78	0.57
Observations	2,033	2,026	1,999	2,030	2,027	2,027

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for a different IV regressions, following the specification described in Section 4.1. *Manager-Salary* is the posterior belief about manager salary, and *Peer-Salary* is the posterior belief about the average peer salary. All the dependent variables correspond to survey questions asked after the elicitation of the posterior beliefs. *E[Future Salary]* corresponds to the expected salary one and five years in the future. *Satisfaction with Pay* and *Satisfaction with Job* are measures in a 5-point scale from very dissatisfied (1) to very satisfied (5). *Ineq. Tol.* measures tolerance for pay inequality in a 3-point scale. 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 4: Effects of Perceived Manager Salary by Distance to Manager

	Log(E[Future Salary])		Effort and Performance		
	+1 year (1)	+5 years (2)	log(<i>Hours</i>) (3)	log(<i>Emails</i>) (4)	log(<i>Sales</i>) (5)
Specification 1 (by No. of Promotions):					
<i>Log</i> (Manager-Salary)					
<i>Closer</i> ⁽ⁱ⁾	0.041 (0.030)	0.204*** (0.059)	0.212** (0.099)	0.170*** (0.052)	0.195 (0.131)
<i>Farther</i> ⁽ⁱⁱ⁾	-0.008 (0.033)	0.086 (0.092)	-0.074 (0.093)	0.019 (0.104)	0.033 (0.285)
Specification 2 (by Promotion Prob.):					
<i>Log</i> (Manager-Salary)					
<i>Closer</i> ⁽ⁱⁱⁱ⁾	0.008 (0.036)	0.200*** (0.059)	0.431* (0.226)	0.185*** (0.061)	0.124 (0.144)
<i>Farther</i> ^(iv)	0.057 (0.038)	0.134 (0.096)	-0.016 (0.135)	0.068 (0.062)	0.289 (0.288)
P-value $H_0 : (i)=(ii)$	0.216	0.229	0.040	0.243	0.657
P-value $H_0 : (iii)=(iv)$	0.322	0.560	0.170	0.212	0.691
Observations	2,033	2,026	602	2,060	791

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column shows results from two regressions (one for each model). In Specification 1, *Closer* indicates managerial positions that are 4 or fewer promotions ahead. In Specification 2, *Closer* indicates a probability of reaching the managerial position within five years of 40% or greater. The regressions in columns (1)–(2) follow the same specification used in columns (1)–(2) of Table 3 (see notes therein for more details), and the regressions from columns (3)–(5) follow the same specifications used in columns (1)–(3) of Table 2 (see notes therein for more details). The only difference with the baseline IV regressions of Tables 3 and 2 is the addition of the binary variable *Closer* as control variable as well as its interaction with *Log*(Manager-Salary). $E[\text{Future Salary}]$ correspond to the expected salary one and five years in the future. *Hours* is the daily number of hours worked. *Emails* is the daily number of emails sent. *Sales* is the sales performance index.

Online Appendix (For Online Publication Only)
Cullen and Perez-Truglia, “How Much Does Your Boss Make?”
June 14, 2021

A Survey Instrument

Below we provide screenshots of the translated survey instrument. The business language of the firm is English so both written and spoken English skills are extremely high, allowing us to pilot the survey with many bilingual employees. After drafting the survey in English, a team of two (employees of the bank) translated the survey so that both versions were available on Qualtrics in English and the local language. We then did all pilot testing of the survey in the local language, meaning we piloted the survey with individuals willing to take the translated version and give detailed feedback about clarity and interpretation in English to us so that we could respond directly with further questions and then instructions for translation edits.

Dear colleagues,

You are invited to participate in a survey study conducted by [Researcher Names] from [Bank Name] in collaboration with a group of academic researchers from Harvard University and other universities from the United States. This survey is intended to teach us more about how [Bank Name] employees learn about their workplace, earnings and career prospects. The purpose is to find ways we can improve our communication about salaries and promotion, and to understand your beliefs about your future career with [Bank Name]. This study is aligned with one of three platforms in the five-year strategy of [Bank Name].

This survey should take less than 30 minutes to complete. All the information provided in this survey is 100% truthful. As a token of our appreciation for your participation, you will be able to earn a minimum of \$9.75 and up to \$700, based on your performance in a game included in this survey.

The rewards will be deposited in your payroll account by the end of Q2.

ALL SURVEY RESPONSES ARE COMPLETELY CONFIDENTIAL. Contact the Office of the Chief Economist should any issue arise.

Thank you in advance for your participation!

Sincerely,

Chief Economist, [Bank Name]

I confirm that I am [Respondent Name] and I would like to take part in this study

>>

To get a general picture of the people answering this survey, we need to know a few things about your background.

Where did you grow up?

Recent research on decision making shows that choices are affected by the context in which they are made. Differences in how people feel, in their previous knowledge and experience, and in their environment can influence the choices they make. To help us understand how people make decisions, we are interested in information about you, specifically whether you actually take the time to read the instructions; if you don't, some results may fail to tell us very much about decision making in the real world. To help us confirm that you have read these instructions, please select the "none of the above" option below. Thank you very much.

In this survey, you will be asked to guess the answer to some questions, and will be rewarded according to the accuracy of your answers. Take the following example:

What is the average height of women in this country (in centimeters)? [Note: we will reward you up to \$2.61 for accuracy]

cm

Note the message "we will reward you up to \$2.61 for accuracy." What we mean by that is that we will use a formula with the ACTUAL average height to reward you. The more accurate your answer is, the more money you will get, up to \$2.61.

Go ahead and provide your guess. This is a practice question, so it will not be scored.



Your guess was 150 cm. The truth is 153 cm. As a result, you would have been awarded \$2.20.

- If you had responded exactly the truth (153 cm), you would have been awarded \$2.61.
- If you had responded 10% above or below the truth (138 cm or 168 cm), then you would have been awarded \$1.05.
- If you had responded 20% above or below the truth (123 cm or 183 cm), then you would have been awarded just \$0.05.

This formula was designed by economists. According to this formula, it is in your best interest to respond honestly.

We will reward you for guessing averages, and also for making other types of guesses. Whenever you see that there is a reward for your guess, please remember that it is in your best interest to respond what you truly believe.

>>

Now, we want to ask you some questions related to salaries. In this survey, we always refer to the monthly basic salary: that is, your monthly salary WITHOUT specific allowances, WITHOUT bonus payments and WITHOUT tax & other deductions. This is the salary specified in your contract.

To make sure that you understand this definition, please try to recall your basic salary and report it here, so we can show you how your answer compares to our records. Please be as exact as possible when reporting this amount, using "." for the decimal separator.

What is your current monthly basic salary from March of 2017? [Note: we will reward you up to \$2.61 for accuracy]

There is a discrepancy between the amount that you reported and our records: you reported a monthly basic salary of \$782, while the administrative records from [Bank Name] indicate an amount of \$730.

Remember, our definition of basic salary EXCLUDES specific allowances, EXCLUDES tax & other deductions and corresponds to March of 2017.

Do you agree with the amount of \$730 shown in our records?

- Yes
- No



**For the remainder of the survey, please keep in mind that all salaries correspond to this same definition:
monthly basic salary for March 2017, WITHOUT specific allowances and WITHOUT tax & other deductions.**



Consider the other employees from the bank who work in your same position (Teller) and unit (Branch 10). According to our records, there are around 50 employees in this group.

What is the average monthly basic salary among all employees in your same position and unit as of March 2017?
[Notes: we will reward you up to \$2.61 for accuracy]

In the previous question, you reported to believe that the average monthly basic salary among employees similar to you was \$848 in March 2017. The next question is designed to assess how confident you feel about your response.

With what probability do you think that the real average could fall in each of the following bins? The probabilities must sum up to 100%. [Note: we will reward you up to \$2.61 for accuracy]

Below \$763.20	<input type="text" value="0"/> %
Between \$763.20 and \$827.13	<input type="text" value="0"/> %
Between \$827.13 and \$868.88	<input type="text" value="0"/> %
Between \$868.88 and \$932.81	<input type="text" value="0"/> %
Above \$932.81	<input type="text" value="0"/> %
Total	<input type="text" value="0"/> %



We have taken a random sample of 5 employees who hold your same position (Teller) and work in your same unit (Branch 10), and calculated the average basic salary among them. With the following set of questions, we want to assess how much you would be willing to pay to obtain this information about average salary.

Below you are presented with 5 hypothetical scenarios. In each scenario, you will be given the choice of either seeing the information about average salary OR receiving extra money as part of your reward for responding to the survey.

We will randomly choose 20 survey respondents. If you are one of these 20 lucky respondents, one of the 5 scenarios will be randomly chosen to be implemented. As a result, it is in your best interest to respond honestly to these scenarios.

Please make your hypothetical choices below, and in the next screen you will find out if your responses will be implemented or remain hypothetical.

Scenario 1: Between the next two options, which one would you prefer?

Information about average salary	\$1.30
<input type="radio"/>	<input type="radio"/>

Scenario 2: Between the next two options, which one would you prefer?

Information about average salary	\$6.52
<input type="radio"/>	<input type="radio"/>

Scenario 3: Between the next two options, which one would you prefer?

Information about average salary	\$26.09
<input type="radio"/>	<input type="radio"/>

Scenario 4: Between the next two options, which one would you prefer?

Information about average salary	\$130.46
----------------------------------	----------

Scenario 5: Between the next two options, which one would you prefer?

Information about average salary	\$652.32
----------------------------------	----------

You have NOT been selected among the 20 participants who will have one of their 5 scenarios implemented. As a result, your choices in the 5 scenarios remain hypothetical.

Please go to the next screen to continue with the survey.



Next, a group of individuals participating in this survey will be chosen to receive some information about the average salary in their same position and unit.

Please continue to the next screen to find out if you will be selected to receive this information.



You have been selected to receive the following information.

We have randomly chosen a random sample of 5 employees who work in your same position (Teller) and unit (Branch 10). The following is the average basic salary in this sample of 5 employees as of March of 2017: \$861.

Please take some time to read and understand this information carefully. When you are ready, proceed to the next screen.

>>

We want to give you the opportunity to re-assess your answer to one of the previous questions. This opportunity is given automatically to all survey participants, regardless of their responses.

What is the average monthly basic salary among all employees who work in your same position (Teller) and unit (Branch 10) as of March 2017? [Notes: we will reward you up to \$2.61 for accuracy]

In the previous question, you reported to believe that the average monthly basic salary among employees similar to you was \$913 in March 2017. The next question is designed to assess how confident you feel about your response.

With what probability do you think that the real average could fall in each of the following bins? The probabilities must sum up to 100%. [Note: we will reward you up to \$2.61 for accuracy]

Below \$821	<input type="text" value="0"/> %
Between \$821 and \$891	<input type="text" value="0"/> %
Between \$891 and \$936	<input type="text" value="0"/> %
Between \$936 and \$1,004	<input type="text" value="0"/> %
Above \$1,004	<input type="text" value="0"/> %
Total	<input type="text" value="0"/> %



Now consider the position Teller Supervisor, which is above your current position.

How many times do you think you would need to be promoted to reach that position (or another position in the same level)?

What is the likelihood that you will be promoted to position Teller Supervisor (or another position in the same or higher level) in the next 5 years?



Consider all employees from the bank who hold position Teller Supervisor. What was their average monthly basic salary as of March 2017? [Note: we will reward you up to \$2.61 for accuracy]

0

>>

In the previous question, you reported to believe that the average monthly basic salary among employees in position Teller Supervisor was \$2,609 in March 2017. The next question is designed to assess how confident you feel about your response.

With what probability do you think that the real average could fall in each of the following bins? The probabilities must sum up to 100%. [Note: we will reward you up to \$2.61 for accuracy]

Below \$2,348	<input type="text" value="0"/> %
Between \$2,348 and \$2,544	<input type="text" value="0"/> %
Between \$2,544 and \$2,674	<input type="text" value="0"/> %
Between \$2,674 and \$2,870	<input type="text" value="0"/> %
Above \$2,870	<input type="text" value="0"/> %
Total	<input type="text" value="0"/> %



We have taken a random sample of 5 employees in position Teller Supervisor, and calculated the average basic salary among them. With the following set of questions, we want to assess how much you would be willing to pay to obtain this information about average salary.

Below you are presented with 5 hypothetical scenarios. In each scenario, you will be given the choice of either seeing the information about average salary OR receiving extra money as part of your reward for responding to the survey.

We will randomly choose 20 survey respondents. If you are one of these 20 lucky respondents, one of the 5 scenarios will be randomly chosen to be implemented. As a result, it is in your best interest to respond honestly to these scenarios.

Please make your hypothetical choices below, and in the next screen you will find out if your responses will be implemented or remain hypothetical.

Scenario 1: Between the next two options, which one would you prefer?

Information about average salary	\$1.30
<input type="radio"/>	<input type="radio"/>

Scenario 2: Between the next two options, which one would you prefer?

Information about average salary	\$6.52
<input type="radio"/>	<input type="radio"/>

Scenario 3: Between the next two options, which one would you prefer?

Information about average salary	\$26.09
<input type="radio"/>	<input type="radio"/>

Scenario 4: Between the next two options, which one would you prefer?

Information about average salary	\$130.46
----------------------------------	----------

Scenario 5: Between the next two options, which one would you prefer?

Information about average salary	\$652.32
----------------------------------	----------

>>

You have NOT been selected among the 20 participants who will have one of their 5 scenarios implemented. As a result, your choices in the 5 scenarios remain hypothetical.

Please go to the next screen to continue with the survey.



Next, a group of individuals participating in this survey will be randomly chosen to receive some information about the average salary in position Teller Supervisor.

Please continue to the next screen to find out if you will be selected to receive this information.



You have been selected to receive the following information.

We have randomly chosen a random sample of 5 employees in position Teller Supervisor. The following is the average basic salary in this sample of 5 employees as of March of 2017: \$2,087.

Please take some time to read and understand this information carefully. When you are ready, proceed to the next screen.



We want to give you the opportunity to re-assess your answer to one of the previous questions. This opportunity is given automatically to all survey participants, regardless of their responses.

Consider all employees from the bank who hold position Teller Supervisor. What was their average monthly basic salary as of March 2017? [Note: we will reward you up to \$2.61 for accuracy]

In the previous question, you reported to believe that the average monthly basic salary among employees in position Teller Supervisor was \$2,348 in March 2017. The next question is designed to assess how confident you feel about your response.

With what probability do you think that the real average could fall in each of the following bins? The probabilities must sum up to 100%. [Note: we will reward you up to \$2.61 for accuracy]

Below \$2,113	<input type="text" value="0"/>	%
Between \$2,113 and \$2,289	<input type="text" value="0"/>	%
Between \$2,289 and \$2,407	<input type="text" value="0"/>	%
Between \$2,407 and \$2,583	<input type="text" value="0"/>	%
Above \$2,583	<input type="text" value="0"/>	%
Total	<input type="text" value="0"/>	%



Now, we want to ask you a few questions about your job at [Bank Name].

Recall that as of March of 2017, your monthly basic salary was \$730.

What do you expect your basic salary to be one year later, in March of 2018?

[Note: we will compare your response to our own projection of your future salary, and we will reward you up to \$2.61 if your response is close to our projection]

And what do you expect your basic salary to be five year later, in March of 2022?

[Note: we will compare your response to our own projection of your future salary, and we will reward you up to \$2.61 if your response is close to our projection]

How satisfied are you with your current salary at [Bank Name]?

- Very satisfied
- Somewhat satisfied
- Neither satisfied nor dissatisfied
- Somewhat dissatisfied
- Very dissatisfied

Across the thousands of [Bank Name] employees, salaries vary with the nature of work, education, experience, responsibilities, etc. What do you think of wage differentials in the company today?

- They are too large
- They are adequate
- They are too small

Taking all the aspects of your job into account, how satisfied are you with your current job at [Bank Name]?

- Very satisfied
- Somewhat satisfied
- Neither satisfied nor dissatisfied
- Somewhat dissatisfied
- Very dissatisfied



What percentage of employees in your same position and unit were assigned to each of the following KPI ratings as of year 2016?

The probabilities must sum up to 100%.

[Note: we will reward you up to \$2.61 for accuracy]

A1	<input type="text" value="0"/>	%
A2	<input type="text" value="0"/>	%
A3	<input type="text" value="0"/>	%
B	<input type="text" value="0"/>	%
C	<input type="text" value="0"/>	%
Total	<input type="text" value="0"/>	%

In comparison to others, are you a person who is generally willing to give up something today in order to benefit from that in the future or are you not willing to do so?

Please use a scale from 1 to 10, where a 1 means you are “completely unwilling to give up something today” and a 10 means you are “very willing to give up something today”. You can also use the values in-between to indicate where you fall on the scale.

1 (Completely Unwilling)	2	3	4	5	6	7	8	9	10 (Very Willing)
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

The survey is almost over. Now we want to assess your attitudes towards transparency. Remember that all your responses are confidential.

How often do you talk about salaries with coworkers?

-
-
-
-
-

If the bank shared with you data on the average pay for all positions. Which positions would you be most interested to look at? Please rank the following options from 1 (most interesting) to 4 (least interesting) by moving the boxes upward or downward:

You own position

Positions right above your level

Positions two levels above of your own position

Other positions



Currently at [Bank Name], salaries are confidential information. Please consider the following two hypothetical scenarios.

Scenario A: the bank created a website showing the average salary by position/unit, for all positions within the bank.

Would you be in favor or against the creation of a website like this?

- Strongly in favor
- In favor
- I would not care
- Against
- Strongly against

Scenario B: the bank created a website with the list of names and salaries of all its employees, including your name and your salary. As a result, you could look up the incomes of any other employee, and any employee could look up your own income.

Would you be in favor or against the creation of a website like this?

- Strongly in favor
- In favor
- I would not care
- Against
- Strongly against



Thanks for completing the survey!

Your total reward for this survey will be the sum of three amounts:

†+ A fixed fee of \$6.52.

†+ The total rewards for the accuracy of your responses during the survey.

†+ A surprise amount, picked at random from the range \$3.26-\$14.35.

We will transfer your total reward to your [Bank Name] account after the survey collection is finalized, which may take up to 10 weeks. You do not need to contact us any further -- rest assured that we will notify you by email when the reward is deposited to your account.

Did you have any technical or language-related problems when doing the survey?

- Yes
- No



B Email with Invitation to Survey

Dear [Employee's Full Name],

We would like to invite you to participate in a survey for [Bank's Name]'s employees. It takes less than 30 minutes to complete the survey and, as a token of our appreciation, you will be receiving a monetary reward – the average reward is around \$30.

[Follow this link to take the survey](#)

This survey is conducted by [Bank's Name] in collaboration with researchers from U.S. universities such as Harvard University. It will help us understand how to communicate with our employees.

You were selected at random to receive this invitation, and all your responses will be completely confidential.

If you have any difficulty responding to this survey, please reply to this email or use the following contact points:

[Bank's Contact 1]

[Bank's Contact 2]

[Bank's Contact 3]

If the link does not work, just copy and paste the following URL to your Internet browser: [Survey's URL]

Thank you for your participation. Your contribution will help to make [Bank's Name] a better place.

Sincerely,

Chief Economist, [Bank's Name]

C Additional Details and Results

C.1 Pay Inequality

This section provides some summary statistics related to pay inequality in the firm where the experiment was conducted and compares it to the inequality in other organizations.

We start by measuring the overall within-firm inequality. We compute a measure used in other studies (Song et al., 2019): the ratio of the 10th to 90th percentile of the distribution of base salary is 0.21. This degree of inequality is quite similar to that of medium sized firms in the United States: the ratio is 0.19 for the average firm with 5,000–10,000 employees (Song et al., 2019). The results are similar if we use different criteria, such as the ratio between the 90th percentile earner and the median earner.

Next, we can decompose the inequality by the horizontal and non-horizontal variation. Let $S_{i,p}$ be the salary of employee i in peer group p (i.e., pair position-unit). By construction:

$$S_{i,p} \equiv (S_{i,p} - \bar{S}_p) + \bar{S}_p \quad (\text{C.1})$$

Then, we can compute the variance of both sides of the equation:

$$\text{var}_i(S_{i,p}) = \text{var}_i(S_{i,p}|i \in p) + \text{var}_p(\bar{S}_p) \quad (\text{C.2})$$

The total dispersion in salaries ($\text{var}_i(S_{i,p})$) is the sum of the horizontal dispersion ($\text{var}_i(S_{i,p}|i \in p)$), weighted by the share of employees in each position, plus the non-horizontal dispersion ($\text{var}_p(\bar{S}_p)$) across positions. Then, we can express the share of horizontal inequality (S_h) as:

$$S_h = \frac{\text{var}_i(S_{i,p}|i \in p)}{\text{var}_i(S_{i,p}|i \in p) + \text{var}_p(\bar{S}_p)} \quad (\text{C.3})$$

Using data on the universe of employees in the firm where the experiment was conducted, we find that only 4.5% of the differences in base salary are horizontal. We can compare the contribution of horizontal inequality with the ones reported in other organizations. We start with the organization from Card et al. (2012). We obtained data on the regular pay in 2014 for all at the employees at the different campuses of University of California. We define peer groups as the combination of position title and department (e.g., one peer group could be the Assistant Professors at the Business School of UCLA). In our study, horizontal inequality accounts for 4.5% of the overall inequality in base salary. We find estimates in the same order of magnitude across the different UC campuses, ranging from a minimum of 7.0% in UC Merced to a maximum of 19.4% in UC San Francisco.

We can also provide a rough comparison to Baker et al. (1994). Using data from a large

U.S. firm, they report that dummies for the eight job levels (i.e., a coarse measure of the vertical position within the firm) explain around 70% of the variation in logarithm of salary. This evidence suggests that vertical inequality probably explains the vast majority of the salary differences within their firm. We used data from our firm to replicate that regression from Baker et al. (1994). We estimated a regression of the logarithm of base salary on a set of dummies for the nine paybands. We find that, like in Baker et al. (1994), a coarse measure of vertical rank picks up the vast majority of the pay inequality. In data from our firm, the level dummies explain 84% of the variation in salaries, which is comparable in magnitude to the corresponding 70% reported in Baker et al. (1994).

We can also provide some comparisons of the country context. While making these comparisons, however, we have to be careful because the employees in our sample are by no means a representative sample of the whole country. On the contrary, our sample represents the richest, most educated segment of the country. For example, in the country where the firm is located, less than 10% of the population aged 25 or older had a College degree in 2019. This share is much lower than for the United States, where the U.S. Census Bureau estimates that share at around 36.0% in 2019. However, among the employees who participated in our study, 86% of them had a College degree. In other words, our sample at hand is almost two and a half times more educated than the average American.

With those caveats in mind, the results for the country context are presented in Figure C.1. Each panel presents a histogram for one specific preference measure from the *Global Preferences Study* across the 76 countries covered in the data. The preference measures are: positive reciprocity, negative reciprocity, altruism, trust, risk taking and patience. For reference, each histogram highlights in blue the bar corresponding to the country where the firm is located. All six panels from Figure C.1 show that the country we study is about average for the 76 countries covered in the data.

C.2 More Details about the Measures of Effort and Performance

In this section, we provide more details and descriptive statistics about the measures of effort and performance: the number of hours worked, emails sent and sales performance.

The data on emails sent is available for all 2,060 employees in the subject sample. The number of hours worked and sales performance, however, is only available for different subsamples of employees: the hours worked can only be measured for employees working in one of the two headquarter offices, while the sales performance is only defined for employees who have some kind of sales role. We observe the email data for all employees. As a result, there is by construction a full overlap between the number of emails sent and the other two outcomes. To illustrate the overlap between the availability of data on hours worked and

sales performance, Table C.1 offers a simple cross-tabulation. The columns (*Sales Data*) indicate whether we observe a sales outcome during the 3 months post-treatment. The rows (*Hours Data*) indicate whether we observe an outcome for hours worked during the 3 months post-treatment. There is almost no overlap between the measures of hours worked and sales performance: we observe data on both of these outcomes for only 67 out of the 2,060 employees. The reason for the lack of overlap is straightforward: only a minority of employees who work in the headquarter offices (for whom we observe the hours worked) have a sales role, and thus we cannot measure sales performance for the majority of such employees.

Figure C.2 shows the distribution of each of the three measures of effort and performance (i.e., the exact same variables used as dependent variables in columns (1) through (3) of Table 2). Figure C.2.a corresponds to the number of hours worked, while Figure C.2.b corresponds to the number of emails sent and Figure C.2.c corresponds to the sales performance. These outcomes are defined exactly as in the regression analysis: i.e., the logarithm of the daily averages over the 3 months post-treatment.

Figure C.3 provides binned scatterplots showing the association between these outcomes. If these outcomes are all picking up effort and performance, they are expected to be positively correlated with each other. Before presenting these results, however, we must point out that we report the raw associations, and as such they are subject to a number of caveats. First, the timing of the association between these outcomes does not need to coincide exactly: for example, it is possible that an increase in the number of hours worked generates higher sales (and thus a higher sales performance index) months later rather than immediately. Second, there is quite a bit of measurement error and volatility in these outcomes, which can generate significant attenuation biases when looking at their associations. Third, the raw correlations do not control for any other employee characteristics (e.g., due to the nature of the tasks, some positions may involve more emails than others). Despite all these caveats however, we should at least expect the associations between the pairs of outcomes to have the right sign.

Figure C.3.a corresponds to the association between the number of emails sent and the number of hours worked, based on data for the 602 subjects for whom we observe both of these outcomes. We find a positive association: a 1% increase in the number of emails sent is associated to a 0.127% increase in the number of hours worked, and this relation is statistically significant (p-value=0.019). Figure C.3.b corresponds to the association between the number of emails sent and the sales performance index, based on data for the 791 subjects for whom we observe both of these outcomes. We observe a positive association here too: a 1% increase in the number of emails sent is associated with a 0.076% increase in the sales performance (p-value=0.011). In sum, and as expected, we find a positive and statistically

significant association between the different measures of effort and performance.⁶⁹

C.3 Randomization Balance and Descriptive Statistics

In this section, we present some descriptive statistics about the subject pool.

Of the 3,841 invitations sent out, 2,060 individuals completed the main module of the survey. This final sample excludes some participants based on pre-treatment characteristics. We excluded 23 participants who were randomly assigned to have their choices in the information-shopping scenarios executed and, as a result, their surveys were programmed to be automatically terminated. Among these 23 participants, 52% had been randomly assigned to receive information about peer salary and 52% had been assigned to receive information about manager salary. The final sample also excludes 1 subject who failed the training test for the definition of own salary: this subject reported an own salary that was over twice as large as the true own salary and then responded as disagreeing with our definition of own salary. This subject was assigned to later receive information on peer salary and manager salary. Last, we excluded 14 subjects with the most extreme prior misperceptions about peer salary (misperceptions over 100%) and manager salary (misperceptions over 400%): of these 14 subjects, 1 were assigned to receive feedback about peer salary and 3 were assigned to receive feedback about manager salary.

Table C.2 presents some descriptive statistics. Column (1) corresponds to the entire subject pool. On average, subjects are 29 years old and have been working at the firm for five years. 73% of them are female and 86% have a college or higher degree. We can check whether there is balance in observables across treatment groups. Subjects were cross-randomized to receive information about manager and peer salary, which resulted in four treatment groups. In columns (2) through (5) of Table C.2, we break down the average characteristics by each of the four treatment groups. The last column reports p-values for the null hypothesis that each average characteristic is constant across the four treatment groups. The results show that, as is consistent with successful random assignment, the observable characteristics are balanced across treatments.

Table C.3 presents descriptive statistics for different samples of employees. Column (1) corresponds to the universe of employees, while columns (2)–(5) correspond to different subsets of the sample: columns (2) and (3) provide summary statistics for the sample of individuals who were not invited and were invited to the survey, respectively; columns (4) correspond to the employees who were invited but did not respond to the survey; and column (5) corresponds to the final sample of 2,060 survey respondents. By comparing columns (1) and (5),

⁶⁹Figure C.3 does not include the association between the hours worked and sales performance because, as shown in Table C.1 and discussed above, there is almost no overlap between those two measures.

it follows that our sample is quite representative of the universe of employees. Even though some of the differences in gender, age, education and tenure are statistically significant, they are always economically small (given the large sample size, we have enough statistical power to detect even small differences). For instance, the subject pool is 73% female vs. 71% female in the universe, the mean ages are 29.2 vs. 30.1 years old, the shares of College graduates are 86% vs. 87%, and the mean tenures are 4.99 vs. 5.09 years.

The only non-trivial difference between the subject pool and the universe of employees is with respect to salary: the average salary in the subject pool is 28% lower than in the universe of employees. We can use the results from the rest of the columns to figure out where this difference is coming from. The comparison between columns (4) and (5) shows that this difference is not coming from differential response rates: the average salary of the survey respondents is quite similar (just 7.5% higher) to the average salary of non-respondents. The comparison between columns (2) and (3) shows that, instead, the bulk of the difference in mean salary between the subject pool and the universe of employees is coming from the selection of employees to be invited to the survey. By construction, we did not send the survey invitation to employees in the highest paybands. The high salaries of some of the excluded employees, such as the CEO and senior vice-presidents, can account for the difference in average pay between the subject pool and the universe of employees.

We provide some suggestive evidence on the accuracy of perceptions about the promotion opportunities. The average subject thinks 3.65 promotions are needed to reach the managerial position they are asked about. To construct a benchmark, we leverage the fact that increases in pay grade typically, although not always, indicate a promotion. Thus, one reasonable proxy for the number of promotions required to reach the managerial position is the difference in pay grades between the employee and his or her manager. On average, employees were 4.32 pay grades away from the managerial position. This distance seems consistent with the subjects' perceived need of 3.65 promotions to reach the managerial position. Moreover, the perceived number of promotions needed to reach the managerial position is significantly correlated to the actual number of pay grades separating the employee from the managerial position (correlation coefficient of 0.403, p -value <0.001).

Last, we provide more details about the definition of the post-treatment outcomes. For the minority of employees who leave the company during the relevant time window, we use the average outcome between the survey date and the exit date. For outcomes that are based on monthly data, such as sales, the post-treatment period corresponds to the month when the survey was taken 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).

C.4 Training Module on the Definition of Base Salary

In this section we provide evidence that individuals understood the definition of base salary that we were using when eliciting beliefs about salaries.

In the training module of the survey, after we provided details about the definition of base salary, we used an incentivized question to ask respondents to report their own base salaries. Figure C.4 compares the individual's guesses about their own salaries with respect to their actual salaries. More than 80% of respondents provide a guess of own salary that is within 5% of the truth – moreover, the typical respondent reported their salary exactly up to the last digit. This outcome confirms the anecdotal evidence that base salary is the most salient aspect of compensation. The remaining 20% of employees missed the mark, typically by a large margin. While it is possible that some misperceptions about own salary exist, these large differences are more consistent with a misunderstanding of the definition of salary. For example, some respondents seem to have reported their salary after taxes and other deductions while in fact we were asking about the salary before taxes. Some respondents provided a guess that is less than half the actual base salary. These employees probably reported an alternative definition of salary that is only relevant for tax purposes (the word in the native language happens to sound similar to base salary). To make sure that these misunderstandings do not extend to the rest of the survey, to the employees who did not guess accurately we showed the employees what their base salary is and then explained the definition of the base salary once again. After this second round of training, we asked these employees who had misreported their own salary if they agreed with our measure and 87% of them responded affirmatively.

C.5 Heterogeneity of Misperceptions and Willingness to Pay

In this section, we provide some heterogeneity analysis of the misperceptions and willingness to pay for salary information.

Table C.4 presents evidence on the heterogeneity of the average error, absolute error and willingness to pay for information. The first set of three columns correspond to perceptions about average manager salary, while the second set of three columns correspond to perceptions about average peer salary. The first row of Table C.4 shows the averages over the entire subject pool. The rest of the rows break down these averages by different subgroups: females vs. males, above and below 4 years of tenure, higher vs. lower paybands, sales vs. non-sales roles, and front office vs. back office roles. In each of these breakdowns, we report

the p-value of the test of the null hypothesis that the relevant average is the same across the two sub-groups.

The most important result from this table is that the patterns are qualitatively consistent across all different subgroups: the absolute error of manager salary is always large, and there is always a systematic bias that is negative and large; the absolute error of peer salary and its bias is also similar across subgroups; and the willingness to pay is always in the neighborhood of \$250. This evidence suggests that the results are not driven by any specific group of the population.

Due to the large sample size, we have enough statistical power to detect even small differences, and for that reason many of the differences in magnitude are statistically significant. However, a large majority of the differences are economically small. We mention below some of the exceptions. Regarding the error on manager salary, the most notable difference is between front (-10%) vs. back office (-24%) roles. Regarding the absolute error on manager salary, the most notable difference is again between front (25%) and back office (36%) employees, which arises mechanically from the differences in biases. Regarding the willingness to pay for information on manager salary, the most notable difference is by front office (\$201.90) versus back office (\$164.91) roles. Regarding the error on peer salary, the most notable exception is the difference by tenure: a systematic bias of 0% for employees with less than 4 years of tenure vs. 5% for employees with more than 4 years of tenure. Regarding the absolute error on peer salary, there are no notable differences. Regarding the willingness to pay for peer salary, the most notable difference is by tenure: \$218.33 for higher tenure vs. \$174.79 for lower tenure.

C.6 Robustness Checks for the Willingness to Pay Data

In this section, we discuss some limitations of the BDM elicitation method, and some evidence addressing those concerns.

The first concern is that our estimates of willingness to pay may be sensitive to the elicitation method – in particular, the lists of prices given in the hypothetical scenarios may act as a suggestion for what the employees “should” pay for the information. As a robustness check, we can take advantage of the fact that we measured willingness to pay in a followup study (Cullen and Perez-Truglia, 2018). In that study, we measured willingness to pay for a similar piece of information: the average salary among a random sample of five peers. Instead of using the price-list method, however, we used an open-ended variation (Andersen et al., 2006), in which the respondent bids against the computer. The rules are as follows. The respondent’s bid is compared to a price that is determined by a random number generator. If the respondent’s bid is lower than the price, then the respondent gets

a dollar amount equal to the price. If the bid is higher than the price, then the subject gets the information and no dollar amount. The rules of this mechanism makes it a dominant strategy for respondents to bid exactly their true valuation for the information. The results are presented in Figure C.5. The left panel corresponds to the willingness to pay as measured in this study, while the right panel replicates the left panel but using the data from Cullen and Perez-Truglia (2018). The distributions should not be expected to be identical, because they are based on non-overlapping samples of employees which differ in a number of features. Despite these differences, the distributions of willingness to pay are in the same order of magnitude across the two elicitation methods. Moreover, this finding is consistent with evidence from other studies showing that measures of willingness to pay are not equal but still largely similar across different elicitation methods (Brebner and Sonnemans, 2018).

Last, despite our efforts for making the willingness to pay elicitation easy to understand, measurement error may still exist. In other words, while subjects may seem quite heterogeneous in their bids for information, some of that heterogeneity may simply reflect factors such as misunderstanding how the elicitation mechanism works. However, our main objective is not to study the heterogeneity in information value, but mostly to assess the average value of information. As a result, the sources of measurement error are likely to be averaged out.

C.7 Determinants of Willingness to Pay for Salary Information

We find that some employees are willing to pay large amounts to acquire information about the average salary of their peers or managers. Our favorite interpretation for that finding is that the interest emerges not just due to curiosity, but primarily due to the instrumental value of the information. In this section, we provide some suggestive evidence in favor of this interpretation.

We test two hypotheses. The first hypothesis is related to the willingness to pay for information about the average peer salary. One of the instrumental reasons for wanting to buy information about the average peer salary is that this information can be useful for salary negotiations. The employee may want to use the information about peer salary as a bargaining chip – for example, they may plan to take a picture of this information and use it as an argument for a raise with their manager or Human Resources representative. Even if the employee was not planning on using the information in the meeting directly, he or she may still use it to decide whether it would be worthwhile to ask for a raise in the first place. According to this channel, we would expect employees who are under-paid to benefit the most from this information, and thus to be willing to pay the most for it, relative to employees who are over-paid.

The results for this first hypothesis are presented in columns (1)–(3) of Table C.5. Each

of the columns from this table corresponds to a separate interval regression, and all of these regressions control for tenure and past performance evaluations. In columns (1) through (3) the dependent variable is the willingness to pay for information on peer salary. In column (1), the key independent variable is the employee's perceived relative salary: i.e., the (log) difference between the employee's own salary and his or her perception of the average peer salary. The coefficient on relative salary is negative (-141.361) and statistically significant (p-value=0.070). This coefficient implies that, consistent with the hypothesis at hand, the employees who are under-paid relative to their peers are the ones most willing to buy information about the average peer salary. The magnitude of this effect is economically significant too: relative to an employee who earns 10% above the peer average, an employee who earns 10% below average would be willing to pay an additional \$31 ($= -141 \cdot (e^{0.2} - 1)$) for the information on the average peer salary.

Columns (2) and (3) of Table C.5 provide two robustness checks for the results presented in column (1). According to this channel, what should really matter for the willingness to pay for information is not what the employee's true relative salary is, but what the employee believes this relative salary to be. In other words, even if an employee was being over-paid relative to peers, if that employee thinks that he or she is under-paid, then he or she should still be willing to pay more for information about the peer salary. The results reported in columns (2) and (3) are consistent with this conjecture. The specification from column (2) is identical to that from column (1), except that it includes as a regressor the employee's true relative salary instead of the employee's perceived relative salary. The coefficient on the true relative salary reported in column (2) is close to zero (-16.166) and statistically insignificant (p-value=0.869). In turn, column (3) is identical to columns (1) and (2), except that both variables (the perceived and true relative salaries) are jointly included in the model. As expected, the coefficient on the perceived relative salary in column (3) is negative (-181.120), statistically significant (p-value=0.069) and, if anything, larger in magnitude relative to the corresponding coefficient from column (1); in turn, the coefficient on the true relative salary is statistically insignificant (p-value=0.462).

The second hypothesis is related to the willingness to pay for information about the manager salary. One of the instrumental reasons for wanting to buy information about the average manager salary is that this information can be helpful for career planning. For example, an employee may want to decide whether it is worth working harder to get promoted to the manager's position (or another position on the same level). Alternatively, an employee may be considering whether to seek an outside job offer, or may already be sitting on an outside offer. To decide whether it is worth staying at this firm or not, in addition to the current salary, the employee should consider the salary growth potential at the current firm.

According to this channel, the willingness to pay for information about a managerial position should depend on whether the employee can aspire to attain that position or not. If the manager's position is out of reach for the employee, then the salary at that position should be largely irrelevant for the employee's career planning. On the contrary, if the manager's position is within the employee's reach, then the information about the salaries in that position should be more useful for career planning and thus more valuable. To test this hypothesis, we can exploit the variation in the distance between the managerial position and the respondent's own position. For some respondents we elicited the willingness to pay for information about a managerial position that is a few promotions away (e.g., asking a junior analyst in investment banking about the salaries of the senior analysts), while for other employees we elicited the willingness to pay for information about the average salary in a managerial position that was well above the employee's current position.

The results for the second hypothesis are presented in columns (4) and (5) of Table C.5. In these two columns, the dependent variable is the willingness to pay for information about manager salary. In column (4), the key independent variable is the employee's perceived probability of being promoted within five years to the managerial position they are being offered the opportunity to buy information about. More precisely, *Perceived P(Promoted to Manager Position)* was constructed based on one of the questions included in the survey, and takes the values from 0.05 (0-10% probability) to 0.95 (90-95% probability), in steps of 0.05. Consistent with the hypothesis of instrumental value, the coefficient reported in column (4) is positive (160.140) and statistically significant (p-value<0.001). This coefficient is significant in magnitude too: relative to a managerial position that the employee could not aspire to reach (probability of promotion of 0%), the employee is willing to pay an additional \$160 to learn the average salary of a managerial position that the employee can aspire to reach (probability of 100%). The results from column (4) are of course subject to the caveat of potential omitted variable biases. In particular, one concern is that the variable on the probability of being promoted to the position may be just picking up the effects of the employee's own salary. To address this concern, column (5) includes the employee's own salary (in logs) in the regression jointly with the perceived probability of being promoted to the manager's position. Despite the addition of the new variable, the coefficient on the perceived probability of promotion remains similar both in terms of magnitude (160.140 in column (4) versus 160.403 in column (5)) as well as in terms of statistical significance (p-value<0.001 in columns (4) and (5) both).

Indeed, some simple back-of-the-envelope calculations indicate that the instrumental value of salary information can be quite significant (Stigler, 1962). For example, assume that an employee is considering acquiring information about peer salary for use in her salary

negotiations. If the expectation is that, with 50% probability, the information will help negotiate a one-year 10% raise, then the employee should be willing to pay up to two weeks of her salary for the information. And since employees may plan to use the information for multiple decisions (e.g., whether to switch jobs or positions), the value of the information can add up quite rapidly across the different margins. In addition to these back-of-the-envelope calculations, there are other approaches to estimating the value of information that also suggest this valuation can be significant. For example, Conlon et al. (2018) use a structural model to estimate the value of salary information for unemployed individuals. They estimate that the average U.S. college graduate looking for a job should be willing to pay \$817 to acquire full information about the distribution of future wage offers.

C.8 Further Analysis on Learning

In this section, we provide additional analysis and robustness checks regarding how individuals updated their beliefs based on the information provided to them.

Figure C.6 provides a finer analysis of the results presented in Figures 2.a and 2.c, but with the following two differences: rather than showing the results for the control and treatment groups on top of each other, the results are presented in separate figures for the treatment and control groups; additionally, the figures include a scatterplot of the raw data on top of the binned scatterplot, where each of the smaller markers correspond to a different respondent. Figure C.6.a and C.6.b correspond to learning about the manager salary. Figure C.6.a shows that most of the observations lie on the x-axis, meaning that the majority of the subjects who did not receive information reported as posterior beliefs exactly what they had reported as prior beliefs. Some observations are somewhat above and somewhat below the x-axis. This means that some subjects revised their beliefs even when they did not receive any feedback. This could be due to a number of reasons, but most likely reflects the typical measurement error observed with survey measures of beliefs. Some subjects may be somewhat inattentive or suffer from survey fatigue, so they respond with different figures when asked a second time. Additionally, even if they are attentive, subjects may respond with slightly different figures due to typos (e.g., \$9,900 in the prior belief and \$9,090 in the posterior belief) or rounding (e.g., \$9,900 in the prior belief and \$10,000 in the posterior belief).

Figure C.6.b is identical to Figure C.6.a, except it corresponds to subjects in the treatment group (i.e., those who received information about the manager salary) instead of the control group. These employees fall into three different groups. Most employees lie close to the 45-degree line, implying that after receiving feedback that did not coincide with their prior beliefs, they fully updated their beliefs to the value of the feedback. A minority of the employees lie close to the x-axis, meaning that even though they were shown feedback about

the manager salary, they decided to disregard that feedback completely and just continue to believe in their prior beliefs. A third group of employees lie in between the previous two groups (i.e., between the 45-degree line and the x-axis, meaning that they adjusted their posterior beliefs in the direction of the feedback, but only partially so). There are many reasons for why some employees did not incorporate the feedback at all or incorporated it only partially. According to the Bayesian learning model, employees are expected to disregard the signal partially or completely to the extent that the perceived accuracy of the signal is lower than the perceived accuracy of their own prior beliefs. In the extreme case, if the employee was 100% sure about his or her own prior belief, Bayesian learning would suggest that such employee should disregard the signal completely – i.e., the employee would attribute the difference between the signal and the prior belief entirely to the sampling variation in the signal.⁷⁰ Additionally, some of the lack of updating may be due to inattention – for instance, due to survey fatigue, some employees may not have incorporated the feedback simply because they skipped over the screen where the signal is shown without inspecting it closely.

Last, Figures C.6.c and C.6.d are identical to Figures C.6.a and C.6.b, except they correspond to learning about the average peer salary instead of the average manager salary. The conclusions for the peer salary are the same as the conclusions described above for the manager salary.

In principle, it is possible for the signal of peer salary to affect beliefs about manager salary, or vice-versa. We explore this possibility in Figure C.7. Panels (a) and (b) of Figure C.7 just reproduce panels (a) and (b) from Figure 2. Panel (c) of Figure C.7 is similar to panel (b), only that the dependent variable is the revision about the manager salary: that is, this figure measures whether the signals about the average peer salary had a causal effect on the subsequent beliefs about the average manager salary. The slope is close to zero and statistically insignificant. In other words, the respondents used the signal about peer salary to update beliefs about peer salary but did not use signals about the average peer salary to update beliefs about the average manager salary. For the sake of completeness, we panel (d) of Figure C.7 is similar to panel (b), except the dependent variable is the revision about the peer salary: that is, it measures whether the signal about manager salary had a causal effect on the beliefs about peer salary. Note such updating would be practically impossible, because the posterior beliefs about peer salary were elicited before the provision of the signal about the manager salary. As expected, Figure C.7.d shows that the signal about the manager salary does not have a causal effect on the reported beliefs on peer salary. Due to the timing

⁷⁰Recall that the signal is based on a random sample of 5 employees, so it contains some sampling variation by construction.

of the survey, this does not prove that individuals did not extrapolate from the signal of manager salary to the belief about peer salary. However, this seems unlikely given the above evidence that subjects did not extrapolate from peer salary to manager salary.

We can assess whether certain subgroups of the population were more likely to incorporate the signals in their posterior beliefs. Table C.6 measures the heterogeneity in learning rates by subgroups of the population. Each column corresponds to a different subgroup: column (1) breaks down the sample by gender, column (2) by tenure, column (3) by payband, column (4) by sales role and column (5) by front versus back office. The top and bottom halves of each column correspond to the estimated learning rates for each subgroup – e.g., in column (1), the top half corresponds to the results for females and the bottom half corresponds to the results for males. The table also reports p-values from the test of the null hypothesis that the learning rates are equal between the two subgroups. The results from Table C.6 indicate that all the differences in learning rate are small in magnitude. For instance, the learning rate for peer salary is 0.536 for females and 0.456 for males, with a difference p-value of 0.413; while the learning rate for manager salary is 0.673 for females and 0.748 for males, with a difference p-value of 0.285. Moreover, only one of the ten differences (five for peer salary and five for manager salary) is statistically significant at conventional levels.

C.9 Additional Results on Information Diffusion

In this section, we present some additional results on information diffusion.

To aid in the interpretation of the results reported in Section 3.4, Table C.7 provides some basic descriptive statistics of all the main variables used for the analysis. Next, we provide a series of robustness checks for the experimental results discussed in Section 3.4.

In the baseline econometric model, given by equation (7), the variable I_i^M indicates if the individual received information about manager salary indirectly. In the baseline definition, this variable takes the value zero if the subject received the information directly. The rationale behind this specification is that if the individual already received the information directly, then whatever information her or she can receive indirectly through peers is pretty much redundant. Even though each individual receives a signal based on a different sample of five peers, the sampling variation in the signals is limited. Thus, the marginal informational value from observing a second signal is quite limited. As a robustness check, Table C.8 reproduces the results from Table 1 but allowing I_i^M to take non-zero values even when the respondent received the information directly. The results from the alternative specification of Table C.8 are almost identical to the results from the baseline specification of Table 1. Another aspect of the baseline specification of equation (7) that is worth discussing further

is that the dependent variable is based on posterior beliefs: $M_i^{abs} = \left| \frac{M_i^{post} - M_i^{true}}{M_i^{true}} \right|$. This specification focuses on posterior beliefs to be able to have the effects of direct information assignment as a benchmark. Due to the structure of the survey, the direct assignment can only affect posterior beliefs. However, if we only want to measure the effects of the indirect assignment, we can focus instead on the effects of prior beliefs. Table C.9 reproduces the results from Table 1 but using prior beliefs as dependent variable: $\left| \frac{M_i^{prior} - M_i^{true}}{M_i^{true}} \right|$. Moreover, Table C.9 uses the definition of I_i^M that can take non-zero values even when the respondent received the information directly (because the respondent would not see that direct information until after the elicitation of the prior beliefs). The results from the alternative specification of Table C.9 are almost identical to the results from the baseline specification of Table 1.

In the baseline results discussed in Section 3.4, the dependent variable is the degree of misperceptions, as measured by the absolute value of the difference between perceptions and reality: $M_i^{abs} = \left| \frac{M_i^{post} - M_i^{true}}{M_i^{true}} \right|$. An alternative outcome of interest is the degree of bias: $M_i^{bias} = \frac{M_i^{post} - M_i^{true}}{M_i^{true}}$. The average of this outcome measures the systematic tendency to either under-estimate or over-estimate average salaries. The results from this alternative dependent variable are presented in Table C.10, which is identical to Table 1 except that the dependent variable is the raw difference between perceptions and actual salaries instead of the absolute difference. In columns (1) through (5) from Table C.10, the dependent variable is the bias on the manager salary. The negative constant from column (1) is negative (-0.088) and statistically significant (p-value<0.001). This estimate indicates that, in the absence of feedback, posterior beliefs end up under-estimating the manager’s true average salary on average. The coefficient on *Direct* from column (1) is positive (0.105) and statistically significant (p-value<0.001). Moreover, the magnitude of the coefficient on *Direct* is similar (in absolute value) to the estimate for the constant. This result indicates that the direct feedback fully eliminates the tendency to underestimate manager salary. Most importantly, the coefficients on all of the different variables associated with indirect feedback in columns (2) through (5), such as *Closest Peer*, are close to zero, statistically insignificant and precisely estimated. Moreover, columns (6)–(10) of Table C.10 reproduce the same analysis as columns (1)–(5) but look at peer salary instead of manager salary. Again, we find robust evidence of an absence of information diffusion.

Intuitively, the results from the baseline specification in Section 3.4 suggest that posterior beliefs get more “compressed” around the truth whenever the subject receives feedback directly, but does not get more compressed around the truth when the subject’s peers receive feedback. By means of graphical analysis, we can inspect whether the feedback impacted the distribution of posterior beliefs in other ways. The results are presented in Figure C.8. Fig-

ure C.8.a provides histograms of the percentage difference between posterior beliefs about the average manager salary and the true average, broken down by whether the subject received a signal of the manager salary (*Direct Feedback*, in red) or not (*No Direct Feedback*, in gray). Figure C.8.a shows that, consistent with the results presented in Section 3.4, the beliefs are more “compressed” around the truth for subjects who received a signal of the manager salary relative to those who did not. Indeed, according to the Kolmogorov-Smirnov test, we can reject the null hypothesis that those two distributions are the same (p-value<0.001). Figure C.8.b is similar to Figure C.8.a, but it corresponds to the subsample of subjects who did not receive feedback directly and breaks them down by whether their closest peer received feedback (*Indirect Feedback*) or not (*No Indirect Feedback*). This figure is also consistent with the results from Section 3.4 in that there is no evidence of information diffusion: we cannot reject the null hypothesis that the distribution of beliefs are equal between those who received feedback indirectly and those who did not (p-value=0.228). In turn, Figures C.8.c and C.8.d are similar to Figures C.8.a and C.8.b, except that they correspond to average peer salary instead of average manager salary. The results are again consistent with a lack of information diffusion: there is a significant difference in beliefs between subjects who received feedback directly versus those who did not (p-value<0.001, from Figure C.8.c); but there is no significant difference in beliefs between subjects who received feedback indirectly versus those who did not (p-value=0.340, from Figure C.8.d).

Last, we supplement the experimental test of information diffusion with a non-experimental test. The theoretical and empirical evidence in the literature indicates that, in presence of information diffusion, individuals who are most central in a network are the ones who are best informed (Alatas et al., 2016; Banerjee et al., 2013; Mobius and Rosenblat, 2014). We can test this hypothesis in our data by comparing the misperceptions between individuals who are more and less central in the network. 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 institution and emails received from outsiders. These results are based on eigenvector centrality, but the findings are similar with alternative definitions of centrality.

The relationship between misperceptions (measures as mean absolute error) and network centrality is shown in binned scatterplot form in Figure C.9.a. Contrary to the hypothesis of information diffusion, we do not find that misperceptions about peer or manager salary decrease with network centrality. On the contrary, the slopes are slightly positive for manager salary (0.283, p-value=0.181) and peer salary (0.299, p-value=0.003). To illustrate these magnitudes, we can compare the misperceptions between employees who have below-median and above-median centrality. For perceived manager salary, the mean absolute error is 28.0%

for individuals with below-median centrality and 28.2% for individuals with above-median centrality (difference p-value=0.910); we find similar results if we look at perceptions on peer salary. For peer salary, the misperceptions are 10.7% for employees with above-median centrality and 12.3% for employees with below-median centrality (difference p-value=0.002).

An alternative way of assessing information diffusion is by using self-reported data on whether employees communicate with each other. We use a survey proxy for whether employees engaged in information diffusion by asking directly how frequently they discussed salaries in the previous year. If accurate information flows in the network, individuals who reported to have discussed salaries with coworkers should have lower misperceptions (Alatas et al., 2016). The results are presented in binned scatterplot form in Figure C.9.b. Contrary to the prediction of information diffusion, we find misperceptions to be statistically indistinguishable between employees who discussed salaries in the past and those who did not. The slopes between the misperceptions and frequency of communication are close to zero and statistically insignificant for both peer and manager salary. Again, to illustrate the magnitude of these differences we can compare misperceptions across employees who never discussed salaries and employees who discussed it once or more. The manager misperceptions are 27.8% for individuals who discussed salaries vs. 28.5% individuals who did not discuss salaries (difference p-value=0.49). For peer salary, the corresponding misperceptions are 11.5% for individuals who discussed salaries vs. 11.6% individuals who did not discuss salaries (difference p-value=0.88). This evidence suggests that, even if employees sometimes discuss salaries with coworkers, they may be sharing noisy or misleading information, or they must not be processing it properly.

C.10 Reduced Form and First Stage Results

Table 2 presents the main results from the IV estimator. For reference, Table C.11 presents the results from the corresponding reduced form and first stage regressions. The results from the reduced form regression are qualitatively similar to the results from the Two-Stage Least Squares regression. The only difference between the two sets of coefficients exists in terms of magnitudes. The first stage results show that individuals did not fully incorporate the feedback given to them. The IV estimates simply scales the reduced form estimates to correct for the incomplete reaction to the information.

C.11 Additional Specification Checks

In this section, we provide some additional specification checks to the results presented in Section 4.

To be able to capture proportional effects, the baseline specification used for Table 2 used the logarithmic transformation for the dependent variables. Next, we present results under a different approach that does not rely on using the logarithmic transformation. In this alternative specification, the dependent variable is defined as the ratio between the post-treatment outcome and the pre-treatment outcome. As such, this specification should be able to capture proportional effects without the need for the logarithmic transformation. The downside of the alternative specification is that, unlike the baseline specification, it cannot be readily interpreted as an elasticity.

The results are presented in Table C.12. Columns (1) through (3) of Table C.12 are a reproduction of the baseline results from the top panel of Table 2. The specifications in columns (4) through (6) of Table C.12 are identical to columns (1) through (3) except that, instead of using the logarithmic transformation, we use the alternative specification discussed above. Since the dependent variables are different, we cannot compare the magnitudes of the coefficients between the baseline specification (columns (1)–(3)) and the alternative specification (columns (4)–(6)). However, we should at least expect them to be qualitatively consistent (i.e., in terms of sign and statistical significance). In terms of the sign, the results are highly consistent between the two specifications: the effects of manager salary are always positive, while the effects of peer salary are always negative. In terms of statistical significance, the results are weaker under the alternative specification (3 out of the 6 coefficients are statistically significant at the 5% level) than under the baseline specification (5 out of the 6 coefficients are statistically significant at the 5% level).

The baseline specification used for Table 2 controls for some basic characteristics of the employee: salary, tenure, and five productivity rating dummies. Table C.13 assesses how robust the results are when we include some additional employee characteristics as control variables. Columns (1) through (3) of Table C.13 are a reproduction of the baseline results from the top panel of Table 2. The specifications in columns (4)–(6) of Table C.13 are identical to columns (1)–(3) except for the inclusion of an additional control variable: a binary variable indicating if the employee is female. The inclusion of this additional control variable makes almost no difference: the results reported in columns (4)–(6) are qualitatively and quantitatively similar to the baseline results reported in columns (1)–(3). The specifications in columns (7)–(9) are identical to columns (1)–(3) except for the inclusion of the following additional control variables: a set of 29 indicator variables corresponding to the location where the employee works. Again, the results are quantitatively and qualitatively robust to the inclusion of the additional controls. Last, the specifications in columns (10)–(12) are identical to columns (1)–(3) except for the inclusion of the following additional control variables: a set of 180 indicator variables corresponding to the employee’s position title. Relative

to the number of observations (602, 2,060 and 791 in columns (1), (2) and (3), respectively) 180 is a large number of control variables. Moreover, some of those categories include a single subject, and thus the inclusion of those indicator variables is numerically equivalent to dropping those observations. For these reasons, the results from this demanding specification must be taken with a grain of salt. In any case, the results under this demanding specification are consistent with the baseline specification, although weaker in terms of statistical significance: 3 out of the 6 coefficients are statistically significant at the 5% level in columns (10)–(12), while 5 out of the 6 coefficients are statistically significant at the 5% level in the baseline specification given by columns (1)–(3).

Next, we assess the robustness checks to a couple of different features of the baseline specification. First, in the baseline specification it is important to control for prior gaps in beliefs. One potential concern is that, due to non-linearities, failing to control for prior beliefs flexibly may bias the estimates. This robustness check is presented in Table C.14. Columns (1) through (3) of Table C.14 are a reproduction of the baseline results from the top panel of Table 2. In turn, the specifications in columns (4)–(6) of Table C.14 are identical to columns (1)–(3) except for the addition of flexible controls for the prior gaps in beliefs about the manager salary and the peer salary. More specifically, in addition to including the prior gaps linearly, we also include sets of dummies for the deciles of the prior gaps (nine dummies for manager salary and nine dummies for peer salary). The findings are robust: the results from columns (4)–(6) are similar to the baseline results in columns (1)–(3), both qualitatively and quantitatively.

The second robustness check presented in Table C.14 relates to the role of belief certainty. When individuals receive information, they may not only change the first moment of the posterior belief (i.e., the mean) but may also change other moments of the distribution (e.g., the dispersion). For example, when an individual receives a signal x he or she may not only shift the mean of the probabilistic belief towards x but also make it more compressed around x (i.e., become more certain). This creates a potential concern: our specification may be attributing all the effects on behavior to the shifts in the first moment of beliefs, when in reality some of the effects are due to higher moments (e.g., certainty). It is unlikely that effects on certainty would explain all the effects we document, however. If, for example, the effects on behavior were due to certainty, we would expect the information to affect behavior in the same direction regardless of whether individuals are updating their mean belief up or down. Instead, our evidence suggests that the effects on behavior go in different directions depending on whether the individual is updating the mean belief up or down. It is still possible, however, that there is a certainty channel and is contaminating the estimates to some extent – and for that reason, we address this concern empirically below.

To address these concerns, we leverage the fact that we elicited not only the mean of the posterior belief but also its distribution. Thus, we can include in the IV regressions the higher moments of the posterior beliefs as additional control variables. If the higher moments are responsible for the effects on behavior, then the IV coefficients should get closer to zero, perhaps even all the way. If, as argued above, the higher order beliefs are not doing any of the heavy lifting, the IV coefficients should remain unchanged when we add these additional controls. The results are presented in columns (7)–(9) of Table C.14, which are identical to the baseline specification from columns (1)–(3), except for the addition of eight control variables related to the distribution of posterior beliefs (four variables for the manager salary and four variables for the peer salary). For the manager salary, the four variables are: the perceived probability that the true average manager salary falls between -10% and -2.5% of the posterior belief of the average manager salary; the perceived probability that the true salary falls between -2.5% and +2.5% of the posterior belief; the perceived probability that the true salary falls between +2.5% and +10% of the posterior belief; the perceived probability that the true salary falls above +10% of the posterior belief. And the four variables on peer salary are defined in an analogous way to the four variables on manager salary. Note that by including this set of variables, we are controlling for the higher moments in a flexible way – that is, we do not need to make any functional form assumptions to estimate a certainty parameter. The findings are robust: the results from columns (7)–(9) are similar to the baseline results in columns (1)–(3), both qualitatively and quantitatively.

C.12 Effects on Other Email-Based Outcomes

In this section, we provide further analysis of the effects of salary perceptions on the employee’s email activity.

In Section 4, we use the number of emails sent by the employee as dependent variable. Our email data allows us to construct alternative dependent variables. The results are presented in Table C.15. Column (1) corresponds to the results from the baseline outcome: the number of emails sent. By construction, these coefficients are identical to the coefficients from column (2) of Table 2. In column (2) of Table C.15, the dependent variable is the number of emails received. While the number of emails sent is more directly linked to the employee effort, there is still an indirect effect of effort on emails received. For instance, if an employee increases the number of emails sent, we would expect some of those additional emails to be replied and thus to translate into additional emails received. Additionally, if the employee takes on more responsibility while working longer hours, that could also translate into more emails received. The results suggest that the two coefficients for emails received (column (2)) are qualitatively consistent with the corresponding coefficients for emails sent (column (1)),

only that smaller in magnitude – this difference is statistically significant for the coefficient on manager salary (p-value=0.096), but statistically insignificant for the coefficient on peer salary (p-value=0.310).

The rest of the columns from Table C.15 break down the effects on the emails sent by the identity of the receivers. First, we break down the number of emails sent by emails sent to emails accounts inside the same firm (column (3)) and emails sent outside of the firm (column (4)). The effects on these two outcomes are qualitative consistent – moreover, we cannot reject the null hypothesis that the two effects are equal to each other (p-values of 0.268 and 0.650 for the coefficients on manager and peer salary, respectively). The last three columns of Table C.15 break down the emails sent to other employees by the rank of those receivers: employees who are in higher paybands (column (5)), employees who are in the same payband (column (6)), and employees who are in lower paybands (column (7)). The results suggest that while the effects of manager salary operate mainly through emails sent to same- and higher-ranked employees, the effects of peer salary operate mainly through emails sent to lower-ranked employees. However, these results have to be taken with a grain of salt, because we cannot reject the null hypothesis that the coefficients are the same across the three columns (p-values of 0.161 for manager salary and 0.464 for peer salary).

C.13 Heterogeneity of Effects of Manager and Peer Salary

Table C.16 presents results on the heterogeneity of the effects of manager and peer salary. In all the regressions, the dependent variable is the number of emails sent. We focus on this outcome because it is available for the entire subject pool and thus it provides the most precision to detect heterogeneity. The top and bottom halves of the table present the coefficients estimated for different subgroups of the population. The first five columns use the same splits employed in the other analyses of heterogeneity: female vs. male (column (1)), higher vs. lower tenure (column (2)), higher vs. lower paybands (column (3)), sales vs. non-sales roles (column (4)), and front-office vs. back-office roles (column (5)). The bottom of Table C.16 provides p-values for the test of the null hypothesis that the coefficients are equal across a given pair of subgroups. We do not find any statistically significant evidence of heterogeneity: none of the ten differences are significant at conventional levels. This evidence suggests that our results are not driven by any specific group of the population. However, due to the precision of the coefficients, we cannot reject moderate differences either. The last column (column (6)) presents results for an additional heterogeneity (by the perceived productivity rank of the employee) that is discussed in the section below.

C.14 Linearity of the Effects on Behavior

In the baseline model, we make the functional form assumption that the relationship between salary perceptions and behavior is log-log linear. In this section, we discuss and relax that assumption.

The baseline specification is simplest one, and is the most common specification in the literature on relative income concerns (see e.g., Senik (2004); Luttmer (2005); Clark et al. (2008); Clark and Senik (2010)). To assess whether this is a reasonable approximation, we provide binned scatterplot versions of the IV regressions from Table 2. The results are presented in Figure C.10. The three panels on the left correspond to the coefficients on manager salary and the three panels on the right correspond to the peer coefficients on peer salary. Each row of two panels correspond to a different regression: the first row corresponds to hours worked, the second row corresponds to emails sent, and the third row corresponds to sales performance. The statistical power available to conduct this type of analysis is limited, so one should not conclude the effects are perfectly linear. However, these binned scatterplots suggests that the log-log linear model is a reasonable approximation. Moreover, these binned scatterplots show that the results do not seem to be driven by outliers, in that the linear regression is not driven by any single bin.

C.15 Symmetry of the Effects on Behavior

Another functional form assumption from the baseline model relates to the symmetry of the responses. Let us start with the main object of interest, the vertical comparisons. Our baseline specification assumes that the effects of updating beliefs upwards are the mirror image of updating beliefs downwards. In practice, finding out that the managers are paid more than initially thought may have stronger or weaker effects than finding out that the managers are paid less than initially thought. For example, employees may have more flexibility to adjust their effort upwards in response to good news about the manager pay (e.g. by working extra hours) than to adjust their effort downwards in response to bad news. To allow for this type of asymmetries, we can augment the baseline model given by equation (8) as follows:

$$\begin{aligned} \log(Y_i^{post}) = & \eta_0 + \eta_{mgr}^{up} \cdot \mathbb{1}_{M_i^{post} > M_i^{prior}} \cdot \log(M_i^{post}) + \eta_{mgr}^{down} \cdot \mathbb{1}_{M_i^{post} \leq M_i^{prior}} \cdot \log(M_i^{post}) + \\ & + \eta_{peer}^{up} \cdot \mathbb{1}_{P_i^{post} > P_i^{prior}} \cdot \log(P_i^{post}) + \eta_{peer}^{down} \cdot \mathbb{1}_{P_i^{post} \leq P_i^{prior}} \cdot \log(P_i^{post}) + \xi_i \end{aligned} \quad (C.4)$$

While the baseline model had one parameter to represent the effects of manager salary (η_{mgr}), this extension has two parameters: one parameter for the upward revisions (η_{mgr}^{up}) and

another parameter for the downward revisions (η_{mgr}^{down}). Likewise, the effects of the horizontal comparisons are split in two distinct parameters. It is straightforward to adjust the IV regression given by equations (9)–(11) to accommodate this form of asymmetry, by splitting the endogenous and instrumental variables into pairs of variables (one for upward revisions and another for downward revisions).

Table C.17 presents the regression results from the asymmetric specification. The top panel correspond to the original (symmetric) specification, which are identical to the baseline results reported in Table 2. The bottom panel corresponds to the asymmetric specification discussed above. Under the null hypothesis of symmetric effects, the coefficients η_{mgr}^{up} and η_{mgr}^{down} should be equal (in absolute value) – the bottom of the table reports the p-value of this test.

Regarding the vertical comparisons, we do not find any significant evidence for this type of asymmetry. We cannot reject the null hypothesis of symmetric effects in any of the seven tests reported in this table. This evidence indicates that the symmetric specification used in the baseline model constitutes a reasonable approximation. However, due to power limitations, we cannot rule out small or moderate asymmetries either. More precisely, the coefficients on downward revisions are quite imprecisely estimated (e.g., in column (1), the standard errors for the downward revisions are six times larger than for the upward revisions). This difference in precision comes from the fact that only a minority of employees end up making downward revisions.

For the horizontal comparisons, the baseline model makes two assumptions about symmetry. The first type of asymmetry is equivalent to the one discussed above for the case of vertical comparisons: the effects of updating beliefs upwards are the mirror image of updating beliefs downwards. Regarding the horizontal comparisons, Table C.17 shows that we reject the null hypothesis of symmetric effects in only one of the seven tests (for sales performance, with a p-value of 0.002). For this outcome, it seems like employees are more responsive to bad news (i.e., peers earning more than previously thought) than to good news. However, given the lack of consistency in the direction of the asymmetry across outcomes, this result is most likely spurious. Again, these findings suggest that the symmetric assumption from the baseline model is a reasonable approximation. However, the statistical power available to conduct this type of analysis is limited, so one should not conclude the effects are perfectly symmetric.

The second form of asymmetry that may arise in horizontal comparisons is the following: the effect of being 1 percent below the peer average salary is the mirror image of the effect of being 1 percent above the peer average.⁷¹ Indeed, there is evidence from related studies

⁷¹Since the perceived manager salary is always above the own salary, this form of asymmetry does not

on this type of asymmetry: while retention goes down when individuals are paid less than the average peer, retention does not go up as much when individuals are paid more than the average peer (Card et al., 2012; Dube et al., 2019; Breza et al., 2018). To allow for this type of asymmetry in horizontal comparisons, we can augment the baseline model given by equation (8) as follows:

$$\log(Y_i^{post}) = \eta_0 + \eta_{peer}^{below} \cdot \mathbb{1}_{P_i^{post} \leq O_i} \cdot \log(P_i^{post}) + \eta_{peer}^{above} \cdot \mathbb{1}_{P_i^{post} > O_i} \cdot \log(P_i^{post}) + \xi_i \quad (\text{C.5})$$

The variable O_i represents i 's own salary. While the baseline model had one parameter to represent the effects of peer salary (η_{peer}), this extension has two parameters: one parameter for when the own salary is below the average peer salary (η_{peer}^{below}) and another parameter for when the own salary is above the average peer salary (η_{peer}^{above}). For example, $\{\eta_{peer}^{below} < 0, \eta_{peer}^{above} = 0\}$ would suggest that employees care about the average peer salary but, once their own salary surpasses the peer average, they no longer care about it. We can expand the IV regression given by equations (9)–(11) to accommodate this form of asymmetry, by splitting the endogenous and instrumental variables for peer salary into pairs (one for when the salary is below the peer average and another for when the salary is above the peer average).

The results for this second form of asymmetry are presented in Table C.18. The top panel presents the results from the original specification, which are identical to the results from Table 2. The bottom panel corresponds to the asymmetric specification. In the null hypothesis of symmetric effects, the coefficients η_{peer}^{below} and η_{peer}^{above} are equal to each other – the bottom of the table reports the p-value of this difference test.

We start by discussing the effects on the retention outcome, which is the the form of behavior for which there is evidence of asymmetric responses (Card et al., 2012; Dube et al., 2019; Breza et al., 2018). These results are presented in column (4). For this outcome, we find a strong asymmetry and in the same direction reported in prior studies (Card et al., 2012; Dube et al., 2019; Breza et al., 2018). The coefficient on peer salary large (0.489) and highly significant when the own salary is below the peer salary, but close to zero (-0.099) and statistically insignificant when the own salary is above the peer salary. Most important, the difference between these two coefficients is statistically significant (p-value=0.047).

On the other hand, we do not find robust evidence for this type of asymmetry in the other outcomes. If anything, the point estimates are skewed in the opposite direction, with two of the differences being statistically insignificant (p-values 0.518 and 0.974) and one being borderline statistically significant (p-value=0.078). Once again, the statistical power

apply to vertical comparisons.

available to conduct this type of analysis is limited, so one should not conclude the effects are perfectly symmetric. However, the evidence does indicate that the symmetry assumption from the baseline model is a reasonable approximation – with the exemption of the retention outcome, for which the symmetric model masks substantial asymmetry.

C.16 Potential Mechanisms Behind Horizontal Comparisons

In Section 4.5 we presented suggestive evidence that social preferences play a role in horizontal comparisons. However, the evidence does not rule out other explanations. In this section, we discuss and provide tests for some alternative channels.

One possibility is that employees reacted to the average salary of their peers because they used that information to learn about other aspects of their jobs. Being paid less than your peers may be a sign that you are less productive than your peers, and being paid more than your peers may be a sign that you are being more productive. Employees may also use the peer salary to infer what their managers think of them. Being paid less than your peers may signal that your manager thinks less of you, while being paid more than your peers may signal that your manager thinks highly of you. Under the right set of assumptions, these inferences could explain why employees work less hard when they receive feedback about a higher-than-expected peer salary.

We have two survey questions that can help us probe those mechanisms. The first question elicits the employee’s self-perception about his or her own relative productivity, using an incentivized method. At the end of every year, each employee is given a productivity rating on a 5-point scale. We elicited the individual’s perception about the share of employees who received 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. This outcome, which is based on an incentivized question, can take values from 0 (least productive) to 1 (most productive). A higher perceived rank may reflect that the employee thinks more highly of herself. Additionally, since the manager provides key input in the performance review, a higher perceived rank may reflect that the manager thinks more highly of the employee. The second question elicits the employee’s probability of being promoted to the manager’s position within the next five years. The employee should expect a higher probability if she thinks more highly of herself, or if she thinks her manager thinks highly of her.

The effects of perceived salaries on these two survey outcomes are presented in Table C.19. This table uses the same IV specification from Table 3, only that it focuses on two different survey outcomes. Column (1) corresponds to the perceived productivity rank. The average of this outcome, 0.47, indicates that individual’s perceptions about their productivity rank

are accurate on average. In other words, employees do not seem to be systematically over- or under-confident. The coefficient on peer salary is close to zero (0.044), precisely estimated, and statistically insignificant. A 10% increase in peer salary, if anything, has a slight positive effect on the perceived productivity rank of just 0.44 pp. In turn, column (2) of Table C.19 presents the effects on the perceived probability of promotion. This dependent variable ranges from 1 (0%–10%) to 10 (90%–100%). The coefficient on peer salary is close to zero (-0.140), precisely estimated, and statistically insignificant. A 10% increase in peer salary reduces the perceived probability of promotion but by just 0.14 pp. In summary, the evidence from Table C.19 goes against the hypothesis that the individuals reacted to the information on peer salary because they inferred something about their productivity or the opinion of their managers.⁷²

The effect of peer salary 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. This mechanism provides a straightforward explanation for the effects on employee retention. Regarding the effects on effort, however, the connection is less clear. On the one hand, employees may work less hard because they expect to change firms and thus no longer care about their internal reputation. On the other hand, there are reasons why employees could want to work harder. For example, employees may want to work harder to be in a better position to ask for a raise, to ask the firm to respond to an outside offer, or to obtain a positive recommendation from their current employer.

Another potential explanation for the effects of peer salary is that individuals use those perceptions to form beliefs about the returns to effort. This channel predicts heterogeneous effects: when receiving a signal that their peers are being paid more than expected, the least productive individuals should infer that the returns on effort are higher and should therefore work harder; on the contrary, the most productive individuals should infer that the returns on effort are lower and should then work less hard. We can provide a test of this channel using the survey data on perceived productivity rank. The results are presented in the last column of Table C.16, where we break down the effects of peer salary by whether the employee's perceived productivity rank is below or above the median. This channel predicts that the effects of higher peer salary should be positive for individuals with below-median perceived productivity and negative for the rest. We do not find any evidence of this form of heterogeneity: the effects of peer salary in these two groups are close to each other and their

⁷²For the sake of completeness, column (1) of Table C.19 also reports the coefficient on manager salary. This coefficient is also close to zero (0.000), precisely estimated, and statistically insignificant. A 10% increase in manager salary increases perceived productivity rank by less than 0.01 pp. This constitutes evidence that employees did not react to the manager salary because they inferred something about their productivity or their manager's opinion. In column (2) we cannot estimate the effect of manager salary because the information on manager salary was provided after the elicitation of the perceived productivity rank.

difference is statistically insignificant ($p\text{-value}=0.788$). This constitutes evidence that the effects of horizontal comparisons do not operate through beliefs about the returns to effort.

Last, it is possible that employees do not care about how much they are paid relative to their peers but are averse to salary inequality. 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 a small or 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 inequality. In other words, while inequality aversion may be important, we would need a different experiment to estimate it.

C.17 Determinants of Horizontal Salary Differences

In this section, we discuss some evidence regarding how horizontal salary differentials are determined and discuss how these findings relate to the interpretation of the experimental results.

We start by providing a bit more detail about the institutional context. Several factors can influence whether one employee gets paid more than a peer. In a nutshell, the salary of a given employee is determined by the Human Resources division in conjunction with the employee's manager. For the sake of simplicity, we start with the case of new hires. The job description determines the pay grade for the position, which is set as part of the bank's overall strategy before a particular individual is assigned the job. This pay grade sets some bounds on the maximum and minimum pay for the position, but those bounds are wide and even overlap across different pay grades. As a result, the pay grade leaves a lot of leeway for horizontal salary differences. Within a given pay grade, the HR division recommends a salary based on market benchmarks that consider not only the position title but also some of the employee's characteristics such as his or her experience. However, that recommendation from HR is not written in stone: the manager has the option to override the HR recommendation. For example, a manager can recommend a salary above the market benchmark if the employee has a competing job offer. For the incumbent employees, the same factors are at play. For example, during the annual review a manager can recommend a higher raise for some employees based on the last year's performance review. And HR has policies that influence the salary growth – e.g., three years of top performance ratings translate automatically into a salary increase based on a fixed formula.

One specific factor that is of particular importance is that of meritocracy. One potential interpretation for the demoralizing effects of horizontal comparisons is that employees want everyone in the same position to get paid the same salary regardless of any differences in

effort or productivity. Another interpretation is that employees do not want to tolerate horizontal salary differences due to non-meritocratic factors such as luck or office politics. The institutional factors described above could support reasons why employee may see the horizontal pay differences as meritocratic (e.g., the raises based on performance reviews) or non-meritocratic (e.g., favoritism due to the manager’s discretion in setting pay). To provide a more quantitative assessment, below we provide quantitative evidence that, indeed, non-meritocratic factors may play a significant role in this context.

We want to assess whether employees end up with higher relative salaries due to meritocratic reasons (e.g., they are more productive or work harder). Let S_i be employee i ’s salary and \bar{S}_{-i} be the average salary among all of i ’s peers. The difference between S_i and \bar{S}_{-i} corresponds to the horizontal salary differential at the time of the experiment. The regression of interest is the following:

$$\log(S_i) - \log(\bar{S}_{-i}) = X_i\beta + \epsilon_i \tag{C.6}$$

Where X_i be a vector of characteristics for employee i such as demographics or measures of past productivity. As a first measure of prior productivity, we use the latest annual performance review. As explained above, the employee’s managers have discretion in determining the employee’s performance rating and the employee’s annual raises. At the end of every year, each employee is given a productivity rating on a 5-point scale. Moreover, the performance reviews and the raises are evaluated around the same time of the year. For these reasons, the performance rating is perhaps the measure of productivity that we should expect to have the strongest association with relative salaries.

The regression results are presented in Table C.20. Column (1) corresponds to a regression where the relative performance review rating is the only explanatory variable. Since we want to explain horizontal salary differences, we calculate the within-peer-group rank in this performance review. Thus, *Performance Rating* takes values from 0 to 1, where 0 would mean that the employee had the lowest performance review rating in the peer group while 1 would mean that the employee had the highest rating. The coefficient on *Performance Rating* is positive (0.052) and highly statistically significant (p-value<0.001). This coefficient is significant in magnitude too: climbing from the bottom to the top of the performance rank in the peer group would be associated with a 5.2% increase in the relative salary (equivalent to 0.33 standard deviations of the dependent variable). While economically significant, this effect is far from explaining all of the horizontal salary differences: the R^2 from column (1) suggests that the past performance ratings can only account for 0.9% of the horizontal differences in salaries.

As complementary evidence, columns (2) through (4) of Table C.20 introduces the three

measures of effort and performance used as the dependent variables in the field experiment: the hours worked, the number of emails sent, and the sales performance. Since we are interested in measuring past performance, we construct the averages of these outcomes in the three months prior to the start of the experiment for each individual. Additionally, we use within-peer-group ranks to facilitate the comparison to the results for the performance review reported in column (1). As a result, the variables *Hours Worked*, *Email Sent* and *Sales* take values from 0 to 1, with 0 corresponding to the bottom of the distribution and 1 to the top of the distribution.

The results from columns (2) through (4) of Table C.20 are roughly consistent with the results from column (1). In column (2), the coefficient on *Hours Worked* is positive (0.016) but smaller in magnitude than the coefficient from column (1) and statistically insignificant (p-value=0.225). In column (3), the coefficient on *Emails Sent* is positive (0.055), similar in magnitude to the coefficient from column (1) and statistically significant (p-value=0.012). In column (4), the coefficient on *Sales* is positive (0.064), similar in magnitude to the coefficient from column (1) and statistically significant (p-value=0.011). On average, the R^2 in columns (2) through (4) is also consistent with the corresponding value from column (1).

In column (5) of Table C.20 we use the employee's demographic characteristics to explain horizontal salary differences. The coefficient on *Female* indicates that female employees are paid 1.3% less than male employees with the same demographic characteristics. These differences can be taken as suggestive evidence that some non-meritocratic factors, such as discrimination, are responsible for some of the horizontal differences – for example, Cullen and Perez-Truglia (2019) provides evidence in this regard using data from this same organization. The variable $\log(\textit{Tenure})$ is meant to proxy for the employee's experience at this specific firm while *Age* is meant to proxy for overall working experience. Both coefficients are positive and statistically significant, indicating that more experienced employees tend to have higher relative salaries. In the case of tenure, the effect could even be mechanical: even if they are not promoted, employees who continue working at this firm would still get raises every year, typically known as cost of living adjustments. To the extent that more experienced employees may be more valuable to the firm, these two coefficients could be taken as evidence of meritocratic pay. The last three variables, *College*, *Business Major* and *Finance Major* are related to the educational credentials of the employees. There is no evidence that, within a peer group, relative salaries are associated with these educational traits.

The R^2 from column (5) of Table C.20 indicates that, taken together, the demographic factors can explain 7.3% of the horizontal salary differences. Moreover, column (6) shows that even when taking the relative performance rating jointly with the demographic characteristics, we can only explain 7.6% of the horizontal salary differences. In sum, the vast

majority of the horizontal salary differences remain largely unexplained. Based on anecdotal accounts, those differences are probably due to sheer luck.⁷³ For example, your initial salary may be higher or lower depending on the representative from Human Resources who was in charge of hiring you, or whether the company was urgently looking to fill a new position at the time they hired you. Moreover, since employees have large misperceptions about peer salary, some employees may end up with lower salaries just because they happen to have more pessimistic beliefs at the time of salary negotiation. Due to the significant role that luck plays, it is plausible that employees perceive the horizontal differences as largely unfair.

C.18 Preferences for Transparency

In this section, we describe two survey questions that elicit more directly whether employees favored or opposed higher transparency. We explained that salaries are currently confidential information at the firm, and asked employees to report whether how they felt about alternative disclosure policies. In the first scenario, we propose the creation of a website showing the same type of information that we provided in our field experiment: i.e., the average salaries by position and unit. Employees could report their support or opposition to this new policy using the following scale: “strongly in favor,” “in favor,” “I would not care,” “against,” or “strongly against”. The results are presented in Figure C.11.a. A majority (65.26%) favors the policy, while 14.22% feel indifferent and only a minority (20.52%) opposes the policy. This survey evidence is consistent with our revealed-preference evidence based on the willingness to pay for salary information, according to which some employees have a lot to gain from having access to more information about salaries. The results are different under the second scenario, in which we offer employees to replace the status quo by a website that shows itemized information about salaries. In other words, you can use this website to look up any specific employee and find out how much they get paid, and other employees can look up your salary. The results, presented in Figure C.11.b indicate that there is little support for non-anonymous disclosure of information: a strong majority (74.83%) opposes the disclosure policy, while 11.84% feel indifferent and only a minority (13.33%) supports it. One plausible interpretation is that while employees value the salary information a lot, they may value their privacy even more.

⁷³These remaining salary differences may still be due to meritocratic factors that we cannot measure. For example, some employees may obtain higher salaries because they are able to attract outside offers (Caldwell and Harmon, 2018) and use them to negotiate raises, which some employees could see as fair. However, we do not have data on outside offers to account for this channel.

C.19 Average Effects from Transparency

In this section, we discuss the average effects of disclosing salary information.

Recall from Section 4.1 that Y_i^{post} is the outcome of interest (e.g., number of emails sent) and T_i^M and T_i^P are the binary variables indicating whether we disclosed information on manager salary and peer salary, respectively. The regression of interest is the following:

$$\log\left(Y_i^{post}\right) = \gamma_0 + \gamma_{mgr} \cdot T_i^M + \gamma_{peer} \cdot T_i^P + X_i \gamma_X + \epsilon_i \quad (\text{C.7})$$

The parameter of interests are γ_{mgr} and γ_{peer} , corresponding to the average effects of disclosing information about the manager salary and about the peer salary, respectively. The vector of additional control variables (X_i) is included to reduce the variance of the error term and thus improve the precision of the estimates. This corresponds to the same set of controls used in the analysis from Section 4.1: the employee’s own salary (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.

Before showing what the average effects of disclosure are, however, it helps to first describe how they relate to the parameters estimated in Section 4. We start by reproducing equation (8), which represents the relationship between salary perceptions and behavior:

$$\log\left(Y_i^{post}\right) = \eta_0 + \eta_{mgr} \cdot \log\left(M_i^{post}\right) + \eta_{peer} \cdot \log\left(P_i^{post}\right) \quad (\text{C.8})$$

In turn, the Bayesian learning model tells us how the disclosure of information affects each of those posterior beliefs:

$$\log\left(M_i^{post}\right) = \log\left(M_i^{prior}\right) + T_i^M \cdot \alpha_{mgr} \cdot \left(\log\left(M_i^{signal}\right) - \log\left(M_i^{prior}\right)\right) \quad (\text{C.9})$$

$$\log\left(P_i^{post}\right) = \log\left(P_i^{prior}\right) + T_i^P \cdot \alpha_{peer} \cdot \left(\log\left(P_i^{signal}\right) - \log\left(P_i^{prior}\right)\right) \quad (\text{C.10})$$

We can combine equations (C.8)–(C.10) to make a prediction about the average treatment effect of disclosing information:

$$\begin{aligned} \log\left(Y_i^{post}\right) = & \eta_0 + \eta_{mgr} \cdot \left(\log\left(M_i^{prior}\right) + T_i^M \cdot \alpha_{mgr} \cdot \left(\log\left(M_i^{signal}\right) - \log\left(M_i^{prior}\right)\right)\right) \\ & + \eta_{peer} \cdot \left(\log\left(P_i^{prior}\right) + T_i^P \cdot \alpha_{peer} \cdot \left(\log\left(P_i^{signal}\right) - \log\left(P_i^{prior}\right)\right)\right) + \nu_i \end{aligned} \quad (\text{C.11})$$

We define $\Delta_{T_i^M} \log\left(Y_i^{post}\right)$ as the effect of disclosing information about the manager salary to individual i on his or her own behavior. That is, the value under $T_i^M = 1$ minus the value

under $T_i^M = 0$:

$$\Delta_{T_i^M} \log(Y_i^{post}) = \eta_{mgr} \cdot \alpha_{mgr} \cdot (\log(M_i^{signal}) - \log(M_i^{prior})) \quad (\text{C.12})$$

Intuitively, the direction of the effect will depend on whether individual i was under- or over-estimating the salary of the manager – and the effect should be zero if the individual’s prior belief was the same as the signal. We can average over individuals to predict the average treatment effect of disclosing the manager salary:

$$\frac{1}{N} \sum_{i=1}^N \Delta_{T_i^M} \log(Y_i^{post}) = \eta_{mgr} \cdot \alpha_{mgr} \cdot \frac{1}{N} \sum_{i=1}^N (\log(M_i^{signal}) - \log(M_i^{prior})) \quad (\text{C.13})$$

In other words, the average treatment effect is the combination of the average bias in prior beliefs multiplied by the degree of belief updating (α_{mgr}) and the degree to which beliefs affect behavior (η_{mgr}). And we can reproduce the same exercise for the horizontal transparency:

$$\frac{1}{N} \sum_{i=1}^N \Delta_{T_i^P} \log(Y_i^{post}) = \eta_{peer} \cdot \alpha_{peer} \cdot \frac{1}{N} \sum_{i=1}^N (\log(P_i^{signal}) - \log(P_i^{prior})) \quad (\text{C.14})$$

We can use the estimates from Section 4 to predict the average treatment effects. The most straightforward prediction is for horizontal transparency: since the prior beliefs are accurate on average, we expect the average treatment effect to be zero. For the vertical transparency, however, we expect a positive average treatment effect, because on average employees underestimate the salary of their managers. However, that average treatment effect should be rather small in magnitude. Take for example the number of hours worked. We estimated an average bias in the prior belief of 0.139 log points, $\hat{\eta}_{mgr} = 0.150$ (column (1) from Table 2) and $\hat{\alpha}_{mgr} = 0.69$ (the slope from Figure 2.b). Thus, we predict that disclosing the manager salary should on average increase hours worked by 1.4% ($= 0.139 \cdot 0.15 \cdot 0.69$). One potential concern with these predictions, however, is that they rely on a number of assumptions such as linearity and symmetry. We can get a more direct measure of the average treatment effects by estimating equation (C.7) directly. And, as a validation exercise, we can compare the results from this simple regression to the predictions discussed above.

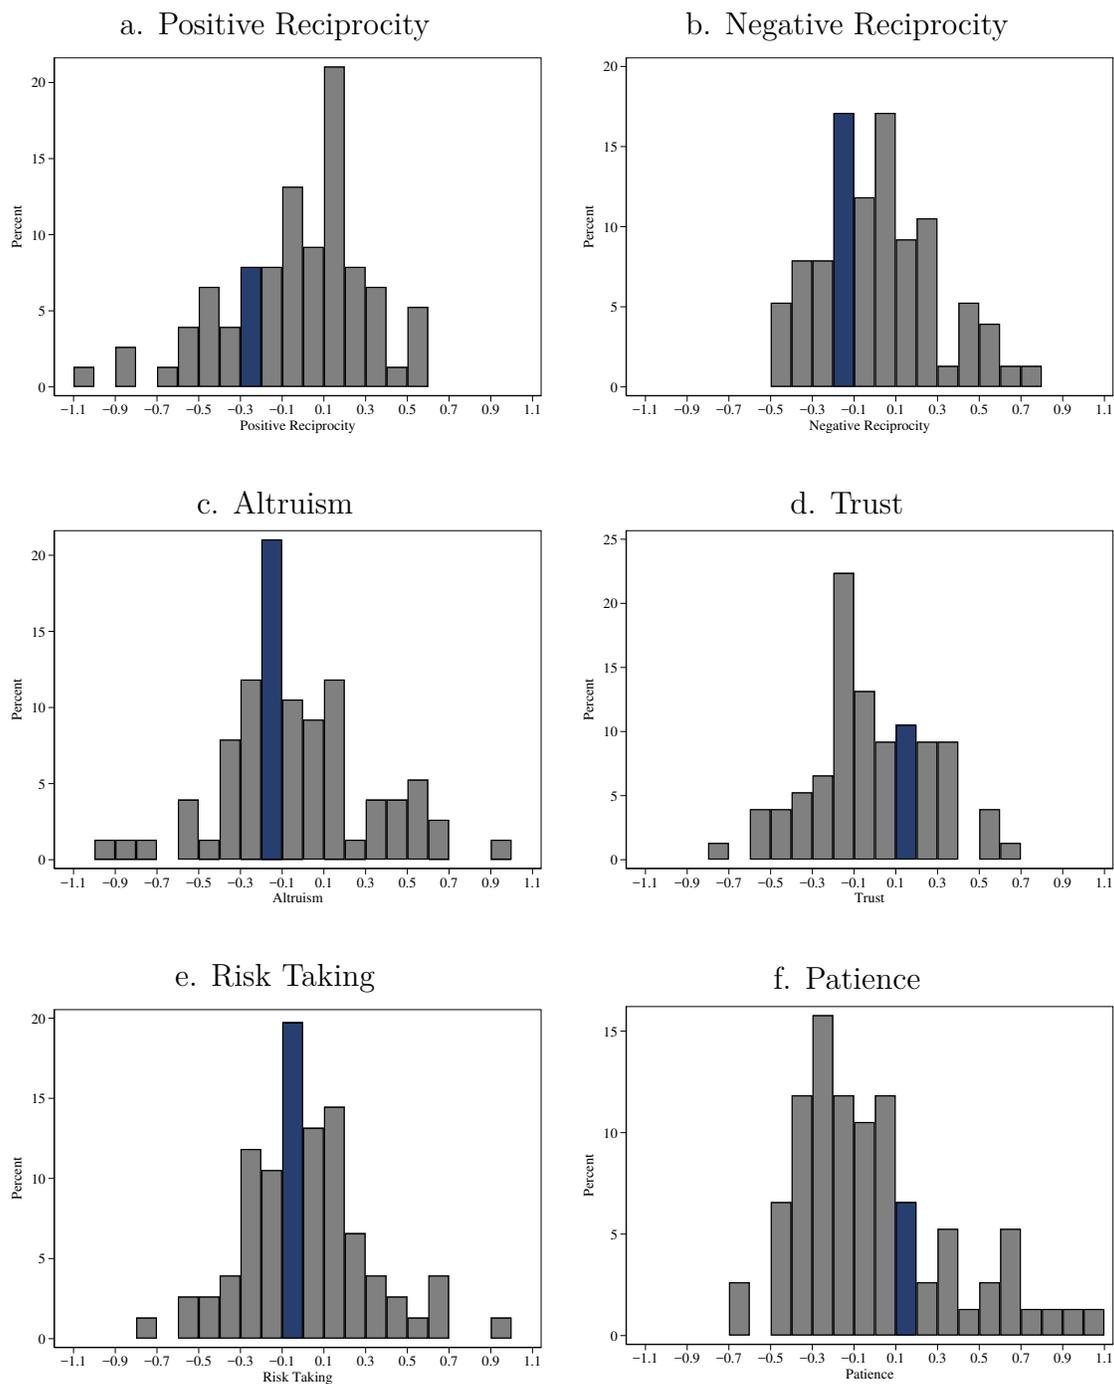
The estimates of equation (C.7) are presented in Table C.21. Different columns correspond to the different measures of effort and performance: the number of hours worked (column (1)), the number of emails sent (column (2)) and the sales performance (column (3)). And consistently with Table 2, in addition to the effects on the post-treatment outcomes, we present the (falsification) coefficients corresponding to the effects on the pre-treatment outcomes. For vertical transparency, columns (1)–(3) provide a consistent picture: on av-

erage, disclosing manager salary has small, positive effects on effort and performance. The effects are 4%, 3.3% and 2.6% for hours worked, number of emails and sales performance, respectively. However, due to lack of precision of the estimates, these coefficients are statistically insignificant. Most important, these estimates are consistent with the predictions discussed above. For example, we predicted that disclosure should increase hours worked by 1.4%, which is statistically indistinguishable from the corresponding effect reported above (4%, from column (1) of Table C.21). The results from Table C.21 are also consistent with the prediction regarding horizontal comparisons that, due to the unbiased beliefs, the average effects from disclosure should be null. Consistent with this prediction, across columns (1) through (3), we find that the average effects of disclosing peer salary are close to zero and statistically insignificant.

Additional References

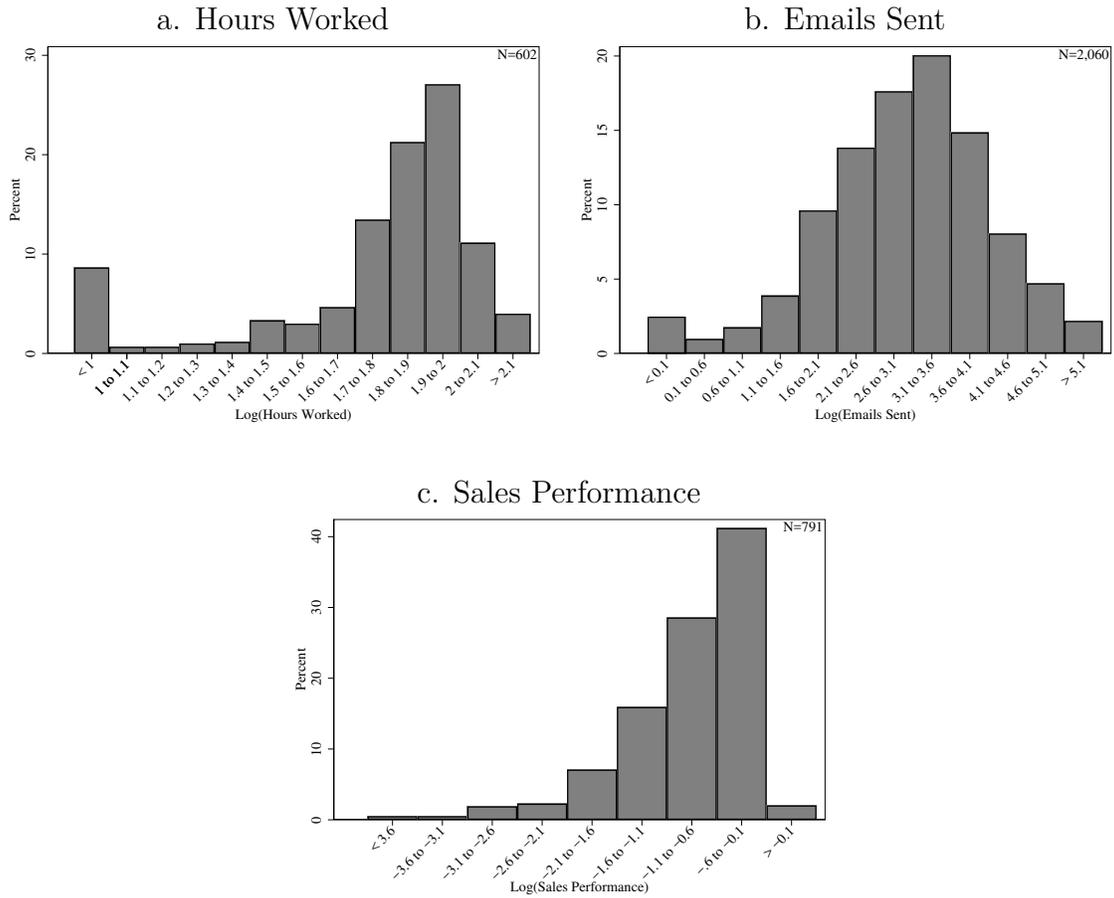
- Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016). Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia. *American Economic Review* 106 (7), 1663–1704.
- Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The Diffusion of Microfinance. *Science* 341 (6144).
- 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.
- Conlon, J., L. Pilossoph, M. Wiswall, and B. Zafar (2018). Labor Market Search With Imperfect Information and Learning. NBER Working Paper No. 24988.
- Luttmer, E. F. P. (2005). Neighbors as Negatives: Relative Earnings and Well-Being. *The Quarterly Journal of Economics* 120 (3), 963–1002.
- Mobius, M. and T. Rosenblat (2014). Social Learning in Economics. *Annual Review of Economics* 6 (1), 827–847.
- Stigler, G. J. (1962). Information in the Labor Market. *Journal of Political Economy* 70(5), 94–105.

Figure C.1: Country Context: Global Preferences Survey



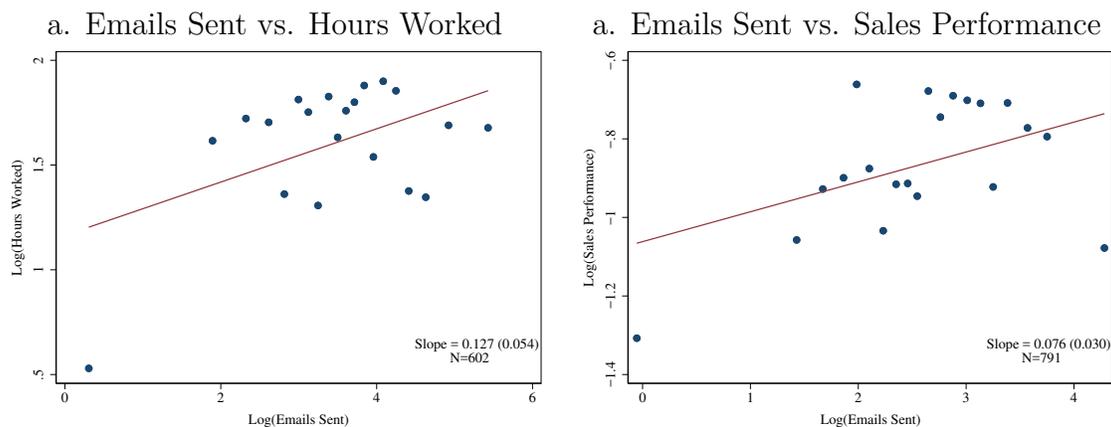
Notes: Distribution of the preference measures across 76 countries. The blue bar corresponds to the country where the firm is located. Data from the Global Preferences Survey (Falk et al., 2018).

Figure C.2: Distribution of Measure of Effort and Performance



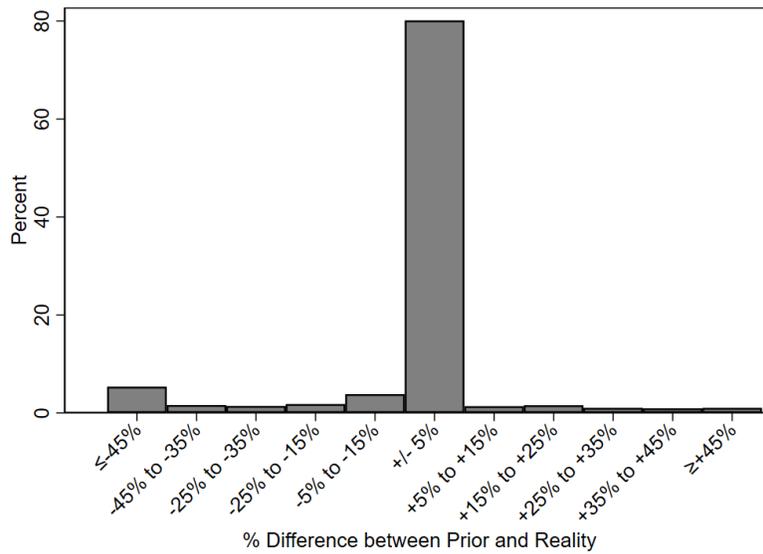
Notes: Histograms of the outcome variables used in columns (1) through (3) of Table 2. Panel (a) corresponds to the (log) mean number of hours worked per day over the 3 months post-treatment. Panel (b) corresponds to the (log) mean number of emails sent per day over the 3 months post-treatment. Panel (c) corresponds to the (log) mean sales performance index over the 3 months post-treatment.

Figure C.3: Association between Different Measures of Effort and Performance



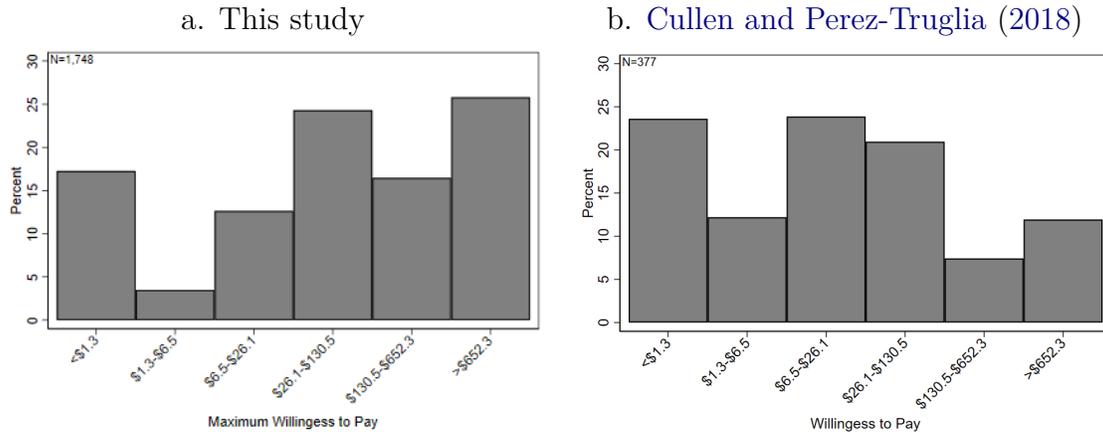
Notes: Binned scatterplots of the association between the outcome variables used in columns (1) through (3) of Table 2. Panel (a) corresponds to the association between the (log) mean number of emails sent per day over the 3 months post-treatment and the (log) mean number of hours worked per day over the 3 months post-treatment. Panel (b) corresponds to the association between the (log) mean number of emails sent per day over the 3 months post-treatment and the (log) mean sales performance index over the 3 months post-treatment. We do not report the association between the hours worked and sales performance because there is almost no overlap between those two measures (see Table C.1). The slopes were estimated with linear regressions, with robust standard errors reported in parentheses.

Figure C.4: Training Module: Eliciting Own Salary



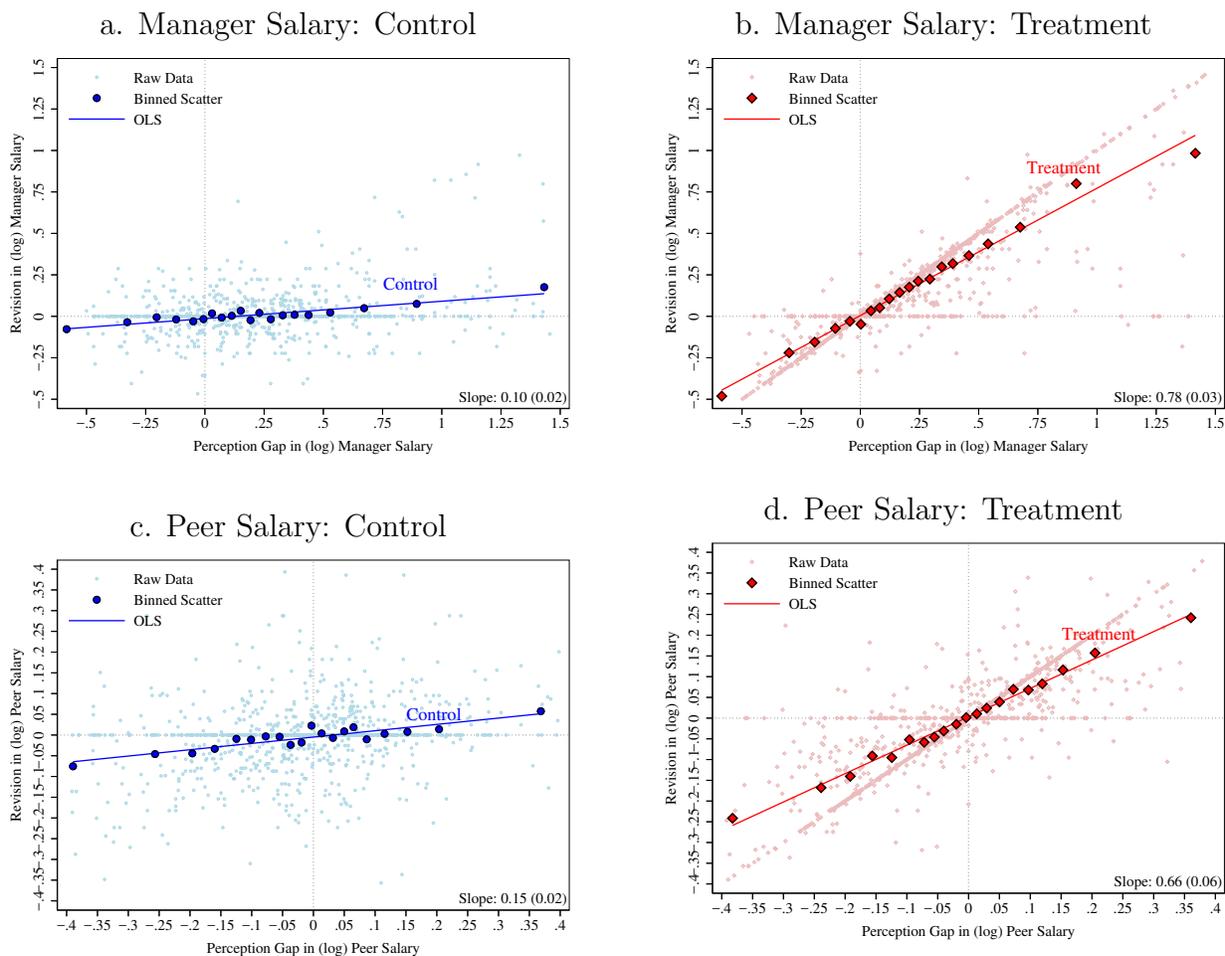
Notes: N=2,060. In the training module, we used an incentivized question to ask employees to report their own base salary. This graph shows the different between the employee's guess and the actual salary (according to the firm's administrative records), divided by the actual salary.

Figure C.5: Willingness to Pay for Salary Information: Comparison to Cullen and Perez-Truglia (2018)



Notes: Panel (a) shows the distribution of the willingness to pay for information about the average salary among peers, using the multiple price-list menu method. The sample is restricted to the subset of respondents with consistent responses across the five price scenarios. Panel (b) is based on data from a follow-up study (Cullen and Perez-Truglia, 2018). It shows the distribution of the willingness to pay for information about the average salary among a sample of five peers, as measured by the respondent’s incentive-compatible bid using the open-ended method. Study participants are a non-overlapping representative sample from the same institution.

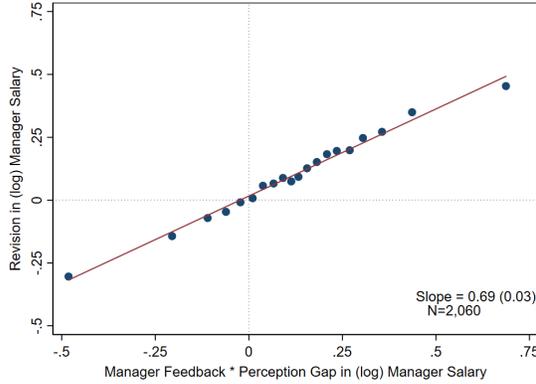
Figure C.6: Effects of Information Provision Experiment on Salary Perceptions: Raw Data



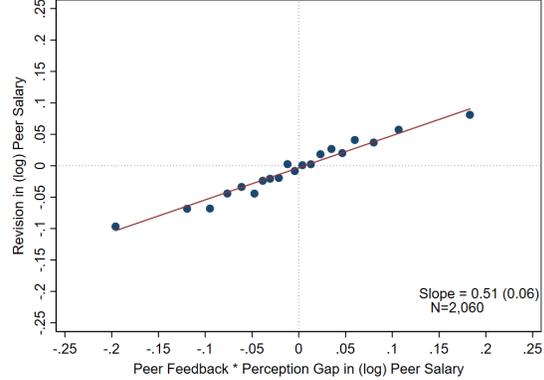
Notes: $N=2,060$. In panel (a), the y-axis is the respondent's update (i.e., the posterior belief about the average manager salary minus the corresponding prior belief) and the x-axis corresponds to the difference between the feedback chosen for the employee (the average salary among the random sample of 5 managers) and the respondent's prior belief. Panel (a) corresponds to subjects who were shown feedback about the manager salary, while panel (b) is identical to panel (a) only that it corresponds to subjects who were not selected to receive feedback about the manager salary. The *raw data* corresponds to a regular scatterplot, where each circle/diamond corresponds to a different respondent (for ease of exposition, we do not plot datapoints outside the range $[-0.5, 1.5]$). The larger circles/diamonds labeled *binned scatter* correspond to a binned scatterplot based on the same data. Panel (c) and (d) are equivalent to panel (a) and (b) except that they are about peer salary instead of manager salary.

Figure C.7: Bayesian Learning in the Information-Provision Experiment

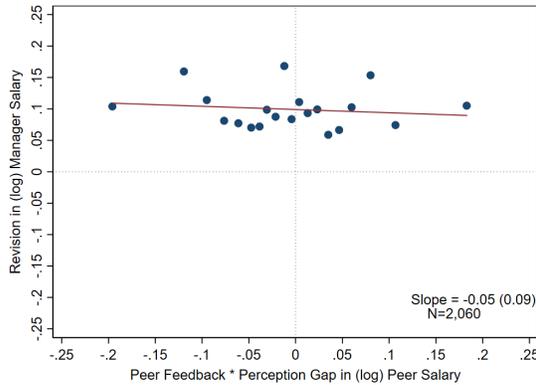
a. Effect of Manager Feedback on Belief About Managers



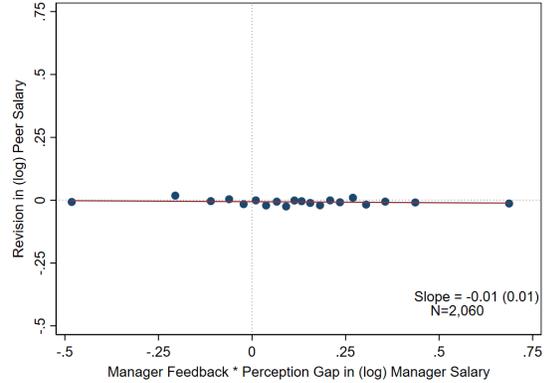
b. Effect of Peer Feedback on Belief About Peers



c. Effect of Peer Feedback on Belief About Managers

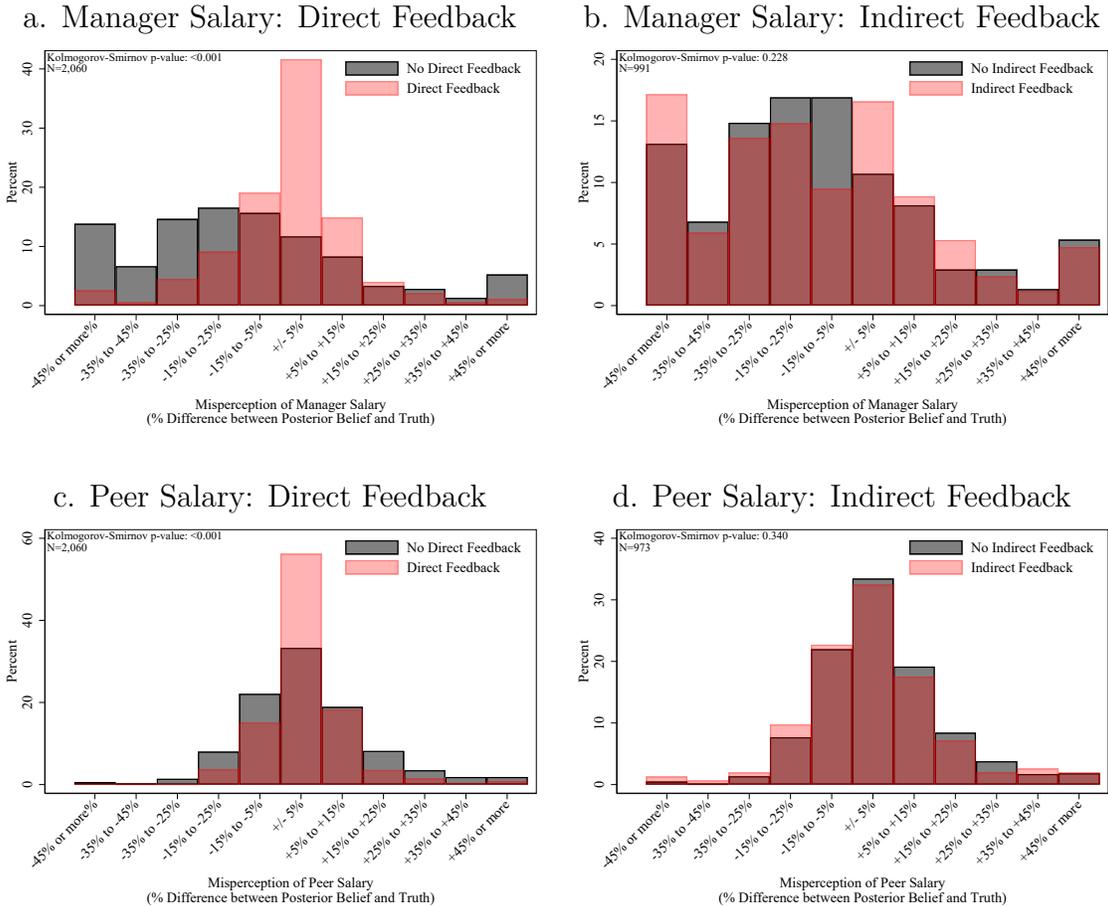


d. Effect of Manager Feedback on Belief About Peers



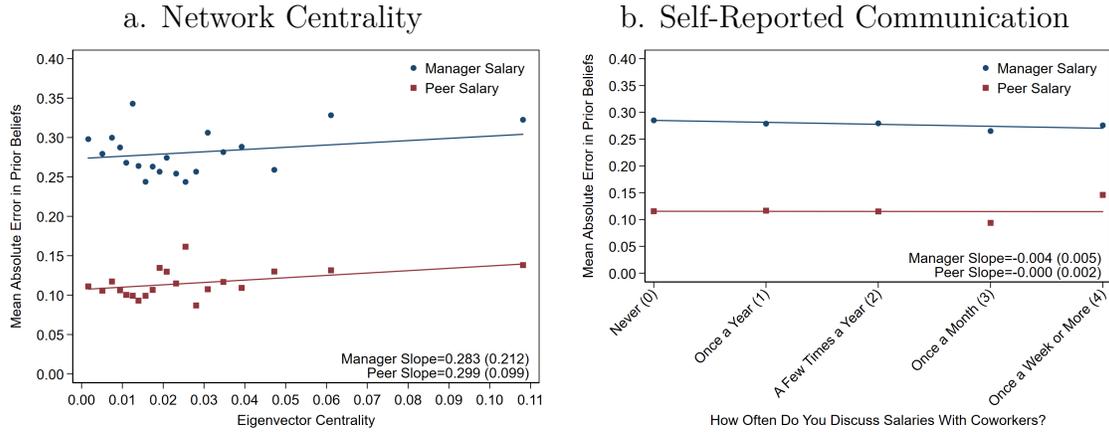
Notes: Panels (a) and (b) present partial regression binned scatterplot of the Bayesian learning equation (6) presented in Section 3.3. The y-axis corresponds to the respondent's update: i.e., the posterior belief minus the prior belief. The x-axis corresponds to the information treatment: the difference between the feedback chosen for the employee (e.g., the average salary among the random sample of 5 managers) and the employee's prior belief, multiplied by a binary 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 was estimated with a linear regression, with standard errors (clustered at position level) reported in parentheses. Panel (a) shows how the feedback about manager salary affects posterior beliefs about manager salary. Panel (b) shows how the feedback about peer salary affects posterior beliefs about peer salary. Panel (c) shows how the feedback about peer salary affects beliefs about manager salary: it is identical to panel (b) except that the dependent variable is the update about manager salary. Panel (d) shows how the feedback about manager salary affects beliefs about peer salary: it is identical to panel (a) except that the dependent variable is the update about peer salary.

Figure C.8: Salary Misperceptions After Direct and Indirect Feedback



Notes: Panels (a) and (b) show the distribution of the difference between the employee’s posterior belief of the average salary (according to an incentivized survey question) and the actual average (according to the firm’s administrative records), divided by the actual salary. Panel (a) breaks down the respondents by whether they were randomly assigned to receive a signal about the average manager salary (*Direct Feedback*) or not (*No Direct Feedback*). Panel (b) is based on the subsample of respondents who did not receive direct feedback about the manager salary, and breaks them down by whether their closest peer received a signal about the average manager salary (*Indirect Feedback*) or not (*No Indirect Feedback*). Panels (c) and (d) are equivalent to panels (a) and (b) but about peer salary instead of manager salary. Panel (c) breaks down the respondents by whether they were randomly assigned to receive a signal about the average peer salary (*Direct Feedback*) or not (*No Direct Feedback*). Panel (d) is based on the subsample of respondents who did not receive direct feedback about the peer salary, and breaks them down by whether their closest peer received a signal about the average peer salary (*Indirect Feedback*) or not (*No Indirect Feedback*). In each panels, the p-value corresponds to the two-sample Kolmogorov-Smirnov test for equality of the two distributions shown in the corresponding panel.

Figure C.9: Non-Experimental Tests of Information Diffusion



Notes: This figure presents binned scatterplot relationships between misperceptions and network centrality (panel (a)) and misperceptions and self-reported communication (panel (b)). In both panels, the y-axis corresponds to the respondent's misperceptions: the prior belief of average (manager/peer) salary minus the true average, divided by the true average. In panel (a), the x-axis measures the employee's network centrality, defined as their eigenvector centrality in the directed network of emails over the three months prior to the completion of the first survey and excluding emails directed outside of the institution and emails received from outsiders. In panel (b), the x-axis corresponds to the answer to the survey question "How often do you talk about salaries with coworkers?" The slopes are estimated with linear regressions, with standard errors (clustered at position level) reported in parentheses.

Figure C.10: Binned Scatterplots of the Effects of Salary Perceptions on Effort and Performance

a. Hours Worked: Manager Salary



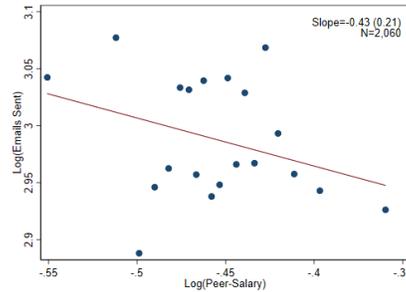
b. Hours Worked: Peer Salary



c. Emails Sent: Manager Salary



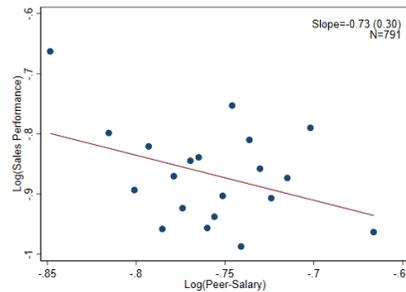
d. Emails Sent: Peer Salary



e. Sales: Manager Salary

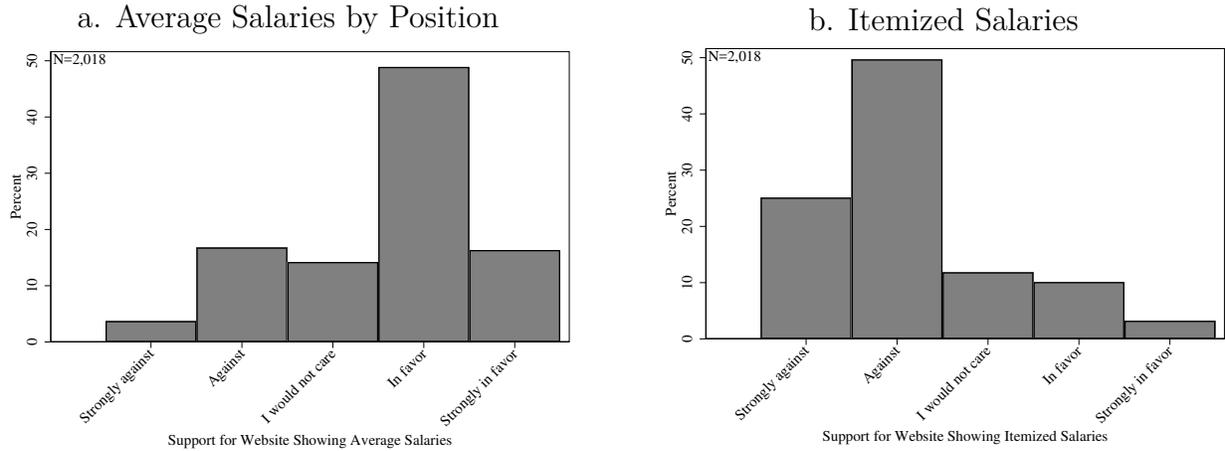


f. Sales: Peer Salary



Notes: This figure presents binned scatterplots corresponding to the regressions reported in Table 2. Each row corresponds to a different IV regression. The dependent variable is 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 sales performance index. The independent variables are: *Manager-Salary* (the posterior belief about the average manager salary) in the first column and *Peer-Salary* (the posterior belief about the average peer salary) in the second column. Each panel reports the corresponding coefficient from the IV regression, with standard errors (clustered at position level) in parentheses. For more details about the regression specification, see notes to Table 2.

Figure C.11: Preferences over Hypothetical Disclosure Policies



Notes: These histograms show the results of two questions included in our survey. Employees are reminded that salaries are confidential information at the firm. In panel (a), employees are asked whether they would support the replacement of the current transparency policy by a website showing the average salary by position/unit for all positions within the bank. In panel (b), employees are asked whether they would support the replacement of the current transparency policy by a website showing the list of names and salaries of all the employees, including your names and salaries.

Table C.1: Availability of Data on Emails Sent, Hours Worked and Sales Performance

Hours Data	Sales Data					
	No		Yes		Total	
	No.	%-Col.	No.	%-Col.	No.	%-Col.
No	721	56.8	737	93.2	1,458	70.8
Yes	548	43.2	54	6.8	602	29.2
Total	1,269	100.0	791	100.0	2,060	100.0

Notes: All 2,060 employees have data available on the number of emails sent. Among these, this table cross-tabulates the availability of data for hours worked and sales performance. The columns (*Sales Data*) indicates whether we observe a sales outcome during the 3 months post-treatment. The rows (*Hours Data*) indicates whether we observe an outcome for hours worked during the 3 months post-treatment.

Table C.2: Descriptive Statistics and Randomization Balance Test

	All	Treatment Group				P-value
	(1)	Manager (2)	Peer (3)	Both (4)	None (5)	
Female	0.73 (0.01)	0.74 (0.02)	0.71 (0.02)	0.74 (0.02)	0.74 (0.02)	0.32
Age	29.20 (0.11)	29.35 (0.22)	29.35 (0.21)	28.92 (0.19)	29.19 (0.22)	0.99
College (or Higher)	0.86 (0.01)	0.84 (0.02)	0.87 (0.01)	0.86 (0.01)	0.86 (0.02)	0.14
Tenure (Years)	4.99 (0.08)	5.14 (0.16)	5.08 (0.15)	4.92 (0.14)	4.79 (0.16)	0.81
Own Salary (Masked)	0.72 (0.01)	0.72 (0.02)	0.72 (0.02)	0.72 (0.03)	0.72 (0.02)	0.93
Avg. Manager Salary (Masked)	2.84 (0.05)	2.80 (0.10)	2.89 (0.10)	2.86 (0.10)	2.80 (0.11)	0.54
Avg. Peer Salary (Masked)	0.72 (0.01)	0.72 (0.02)	0.72 (0.02)	0.73 (0.02)	0.73 (0.02)	0.91
Observations	2,060	510	528	559	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. Manager Salary* and *Avg. Peer Salary* are the true average salaries among the manager and peer groups, respectively. 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 the average manager salary only (column (2)); receiving information about the average peer salary only (column (3)); receiving information about both manager and peer salary (column (4)); and receiving no salary information (column (5)).

Table C.3: Average Characteristics in Subject Pool vs. Universe of Employees

	All	Invited		Responded	
	(1)	No (2)	Yes (3)	No (4)	Yes (5)
Female	0.71 (0.01)	0.65 (0.01)	0.76 (0.01)	0.79 (0.01)	0.73 (0.01)
Age	30.14 (0.06)	31.33 (0.10)	29.13 (0.08)	29.04 (0.12)	29.20 (0.11)
College (or Higher)	0.87 (0.00)	0.92 (0.00)	0.83 (0.01)	0.81 (0.01)	0.86 (0.01)
Tenure (Years)	5.09 (0.04)	5.32 (0.07)	4.90 (0.05)	4.80 (0.08)	4.99 (0.08)
Own Salary (Masked)	1.00 (0.01)	1.39 (0.03)	0.66 (0.01)	0.60 (0.01)	0.72 (0.01)
Avg. Manager Salary (Masked)	3.38 (0.04)	4.61 (0.08)	2.52 (0.03)	2.15 (0.04)	2.84 (0.05)
Avg. Peer Salary (Masked)	1.00 (0.01)	1.39 (0.03)	0.67 (0.01)	0.60 (0.01)	0.72 (0.01)
Observations	(Masked)	(Masked)	3,841	1,781	2,060

Notes: This table presents 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. Manager Salary* and *Avg. Peer Salary* are the true average salaries among the manager and peer groups, respectively. Column (1) corresponds to the entire subject pool. Columns (2) and (3) split the universe of employees by whether they were invited (or not) to participate in the survey. Columns (4) and (5) split the sample of employees invited to the survey by whether they responded to the survey or not. Due to the sensitive nature of the data, we do not report the unit of measurement for the salary variables or the total number of employees in the organization.

Table C.4: Heterogeneity in Misperceptions and Willingness to Pay for Information

	Average Manager Salary			Average Peer Salary			Observations
	Error	Abs. Error	WTP	Error	Abs. Error	WTP	
All	-0.14 (0.008)	0.28 (0.005)	190.74 (6.652)	0.03 (0.003)	0.12 (0.002)	197.01 (6.506)	2,060
By Gender:							
Male	-0.14 (0.016)	0.29 (0.012)	208.03 (12.775)	0.02 (0.007)	0.12 (0.005)	216.40 (12.481)	554
Female	-0.14 (0.008)	0.28 (0.006)	183.77 (7.778)	0.03 (0.004)	0.11 (0.003)	189.22 (7.609)	1,506
Diff p-value	0.64	0.14	0.09	0.32	0.16	0.05	
By Tenure:							
> 4ys	-0.16 (0.010)	0.28 (0.007)	192.01 (9.349)	0.05 (0.005)	0.12 (0.004)	218.33 (9.478)	1,054
≤ 4ys	-0.12 (0.011)	0.28 (0.008)	189.40 (9.471)	0.00 (0.005)	0.11 (0.003)	174.79 (8.828)	1,006
Diff p-value	<0.01	0.86	0.84	<0.01	<0.01	<0.01	
By Payband:							
<i>HigherPayband</i>	-0.16 (0.011)	0.27 (0.008)	187.37 (9.873)	0.03 (0.006)	0.12 (0.004)	206.07 (9.788)	898
<i>LowerPayband</i>	-0.13 (0.011)	0.29 (0.007)	193.44 (9.003)	0.02 (0.004)	0.11 (0.003)	189.66 (8.704)	1,162
Diff p-value	0.01	0.25	0.65	0.38	<0.01	0.20	
By Sales Role:							
<i>Sales</i>	-0.09 (0.011)	0.25 (0.007)	203.12 (9.999)	0.04 (0.005)	0.11 (0.004)	190.83 (9.483)	972
<i>Non – Sales</i>	-0.19 (0.011)	0.31 (0.007)	180.06 (8.884)	0.02 (0.005)	0.12 (0.003)	202.31 (8.941)	1,088
Diff p-value	<0.01	<0.01	0.08	<0.01	0.75	0.37	
By Role:							
Front Office	-0.10 (0.008)	0.25 (0.006)	201.90 (8.139)	0.03 (0.004)	0.11 (0.003)	196.81 (7.780)	1,454
Back Office	-0.24 (0.015)	0.36 (0.011)	164.91 (11.381)	0.02 (0.007)	0.12 (0.005)	197.48 (11.875)	606
Diff p-value	<0.01	<0.01	0.01	0.22	0.02	0.96	

Notes: This table presents average error, absolute error, and WTP for various groups with standard errors in parentheses (clustered at the position level). *Error* is the gap between the perceived (manager/peer) average salary and the true average salary, divided by the true average salary. *Abs. Error* is the absolute value of *Error*. WTP is the willingness to pay for (manager/peer) information, calculated from interval data using the most conservative approach that focuses on the lower bound of each interval (see Section 3.2 for more details). The row “All” corresponds to the full sample and the other rows correspond to different subsamples. P-values correspond to the test of the null hypothesis that the average is equal across the two subgroups.

Table C.5: Determinants of Willingness to Pay for Salary Information

	WTP Peer Info			WTP Mgr Info	
	(1)	(2)	(3)	(4)	(5)
Log(Own Salary) - Log(Perceived Peer Avg. Salary)	-141.361*		-181.120*		
	(77.889)		(99.761)		
Log(Own Salary) - Log(True Peer Avg. Salary)		-16.166	94.952		
		(98.138)	(129.090)		
Perceived P(Promoted to Manager Position)				160.140***	160.403***
				(34.146)	(33.988)
Log(Own Salary)					12.542
					(32.003)
Observations	1,748	1,748	1,748	1,637	1,637

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for a separate interval regression. All regressions control for tenure (in logs) and dummies for performance evaluations in the previous year. In columns (1)–(3) the dependent variable is the willingness to pay for information on peer salary (restricted to the sample of respondents with consistent responses across the five scenarios). In columns (4)–(5) the dependent variable is the willingness to pay for information on manager salary (restricted to the sample of respondents with consistent responses across the five scenarios). *Log(Own Salary)* is the logarithm of the subject’s own salary. *Log(Perceived Peer Avg. Salary)* is the logarithm of the average peer salary according to the subject’s prior belief reported in the survey. *Log(True Peer Avg. Salary)* is the logarithm of the true average peer salary according to the administrative records. *Perceived P(Promoted to Manager Position)* is the subject’s perceived probability of being promoted to the managerial position according to survey data and can take values from 0.5 to 0.95 in increments of 0.05.

Table C.6: Heterogeneity in Learning Rates

	(1)	(2)	(3)	(4)	(5)
	Female	> 4ys	High-Band	Sales	Front-Office
α_{manager}	0.673*** (0.040)	0.630*** (0.056)	0.714*** (0.045)	0.633*** (0.030)	0.682*** (0.030)
α_{peer}	0.536*** (0.080)	0.516*** (0.093)	0.537*** (0.110)	0.568*** (0.069)	0.550*** (0.059)
Observations	1,506	1,054	898	972	1,454
	Male	$\leq 4ys$	Low-Band	Non-Sales	Back-Office
α_{manager}	0.748*** (0.058)	0.791*** (0.036)	0.677*** (0.049)	0.713*** (0.047)	0.697*** (0.057)
α_{peer}	0.456*** (0.054)	0.508*** (0.049)	0.478*** (0.033)	0.458*** (0.083)	0.415*** (0.128)
Observations	554	1,006	1,162	1,088	606
P-value Diff.:					
Manager	0.285	0.016	0.583	0.156	0.811
Peer	0.413	0.940	0.610	0.306	0.338

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Estimation of Bayesian learning equation (6) presented in Section 3.3 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to Table 2. α_{mgr} and α_{peer} correspond to the learning rates for manager and peer beliefs (i.e., the weight that the individual assigns to the signal relative to the weight assigned to the prior belief), which are estimated from separate regressions. Each column corresponds to a different split of the sample (e.g., females vs. males), and the p-values correspond to the test of the null hypothesis that the learning rates are equal across the two subgroups.

Table C.7: Descriptive Statistics about Information Diffusion Analysis

	Manager Salary			Peer Salary		
	All	Received Own?		All	Received Own?	
		Yes	No		Yes	No
	(1)	(2)	(3)	(4)	(5)	(6)
Misperceptions (Posterior Belief)	0.189 (0.201)	0.113 (0.139)	0.271 (0.225)	0.090 (0.111)	0.069 (0.102)	0.114 (0.115)
Information Assignment						
Direct	0.519 (0.500)	1.000 (0.000)	0.000 (0.000)	0.528 (0.499)	1.000 (0.000)	0.000 (0.000)
Indirect: Closest Peer	0.082 (0.274)	0.000 (0.000)	0.171 (0.376)	0.075 (0.263)	0.000 (0.000)	0.158 (0.365)
Indirect: No. Peers	1.346 (2.651)	0.000 (0.000)	2.797 (3.249)	1.313 (2.471)	0.000 (0.000)	2.779 (2.976)
Indirect: (No. Peers > 0)	0.340 (0.474)	0.000 (0.000)	0.707 (0.455)	0.352 (0.478)	0.000 (0.000)	0.745 (0.436)
Indirect: Share of Peers	0.060 (0.109)	0.000 (0.000)	0.125 (0.128)	0.061 (0.108)	0.000 (0.000)	0.129 (0.125)
Observations	2,060	1,069	991	2,060	1,087	973

Notes: This table presents average values of the key variables used for the information diffusion analysis, with standard deviations in parentheses. Columns (1)-(3) refers to the manager salary. *Misperceptions (Posterior Belief)* is the absolute value of the difference between the posterior belief about average manager salary and the true average, divided by the true average. *Direct* is a binary variable indicating whether the subject received the signal on manager salary directly. *Closest Peer*, *No. Peers*, *(No. Peers > 0)* and *Share of Peers* measure if the subject peer's received the signal on manager salary before the individual started his or her own survey. By definition, these variables take the value 0 if the employee received the information directly in his or her own survey. *Closest Peer* is a binary variable taking the value 1 if the individual's closest peer (defined as the peer with whom the employee exchanges the most number of emails in Jan-Mar 2017) received the information. *No. Peers* denotes the number of employees in the employee's peer group who received the information. *(No. Peers > 0)* is a binary variable indicating if *No. Peers* is positive. *Share of Peers* denotes the share of employees in the employee's peer group who received the information. Columns (4)-(6) show the corresponding results from columns (1)-(3) but for peer salary instead of manager salary.

Table C.8: Information Diffusion Analysis: Alternative Specification for Indirect Assignment

	Misperceptions on Manager Salary					Misperceptions on Peer Salary				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Information Assignment										
Direct	-0.160*** (0.011)	-0.160*** (0.011)	-0.160*** (0.011)	-0.160*** (0.011)	-0.160*** (0.011)	-0.045*** (0.005)	-0.045*** (0.005)	-0.045*** (0.005)	-0.045*** (0.005)	-0.045*** (0.005)
Indirect: Closest Peer		0.004 (0.013)					0.010 (0.008)			
Indirect: No. Peers			0.000 (0.002)					0.001 (0.001)		
Indirect: (No. Peers > 0)				0.012 (0.013)					0.008 (0.007)	
Indirect: Share of Peers					0.024 (0.048)					0.008 (0.020)
Constant	0.276*** (0.019)	0.276*** (0.019)	0.277*** (0.021)	0.275*** (0.019)	0.276*** (0.019)	0.111*** (0.011)	0.111*** (0.011)	0.114*** (0.011)	0.110*** (0.011)	0.111*** (0.011)
Observations	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060

Notes: N= 2,060. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the peer group level. This Table reproduces Table 1, only that the variables *Closest Peer*, *No. Peers*, *(No. Peers > 0)* and *Share of Peers* are not forced to take the value 0 if the employee received the information directly in his or her own survey. See the notes to Table 1 for more details.

Table C.9: Information Diffusion Analysis: Alternative Specification based on Prior Beliefs

	(Prior) Misperceptions on Manager Salary					(Prior) Misperceptions on Peer Salary				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Information Assignment										
Indirect: Closest Peer		0.006 (0.018)					0.005 (0.007)			
Indirect: No. Peers			-0.000 (0.003)					0.002* (0.001)		
Indirect: (No. Peers > 0)				0.015 (0.017)					0.002 (0.006)	
Indirect: Share of Peers					0.078 (0.061)					0.014 (0.021)
Constant	0.293*** (0.027)	0.293*** (0.027)	0.292*** (0.030)	0.291*** (0.027)	0.292*** (0.027)	0.090*** (0.009)	0.090*** (0.009)	0.093*** (0.009)	0.090*** (0.009)	0.089*** (0.009)
Observations	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060

Notes: N= 2,060. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the peer group level. This Table reproduces Table C.8, but using as dependent variable the misperceptions in prior beliefs instead of the misperceptions in posterior beliefs. we do not include the *Direct* binary variable because that information is always provided after the elicitation of prior beliefs. Also, the variables *Closest Peer*, *No. Peers*, *(No. Peers > 0)* and *Share of Peers* are not forced to take the value 0 if the employee received the information directly in his or her own survey. See the notes to Table 1 for more details.

Table C.10: Information Diffusion Analysis: Specification with Alternative Dependent Variables

	Bias on Manager Salary					Bias on Peer Salary				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Information Assignment										
Direct	0.105*** (0.014)	0.106*** (0.015)	0.108*** (0.018)	0.123*** (0.026)	0.108*** (0.019)	-0.004 (0.007)	-0.007 (0.007)	0.001 (0.009)	-0.005 (0.013)	-0.002 (0.009)
Indirect: Closest Peer		0.006 (0.026)					-0.020 (0.016)			
Indirect: No. Peers			0.001 (0.004)					0.002 (0.002)		
Indirect: (No. Peers > 0)				0.025 (0.030)					-0.001 (0.013)	
Indirect: Share of Peers					0.022 (0.104)					0.017 (0.044)
Constant	-0.088*** (0.033)	-0.088*** (0.033)	-0.088*** (0.033)	-0.098*** (0.035)	-0.089*** (0.034)	0.013 (0.014)	0.015 (0.014)	0.012 (0.014)	0.014 (0.015)	0.012 (0.014)
Observations	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060

Notes: N= 2,060. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the peer group level. This table is identical to Table 1 except that the dependent variables are not constructed in absolute values: e.g., in columns (1) through (5) the dependent variable is the difference between the posterior belief of the average manager salary and the true average, divided by the true average. See the notes to Table 1 for more details.

Table C.11: Reduced Form and First Stage Results from the Instrumental Variables Estimator

	Effort and Performance		
	(1) log(<i>Hours</i>)	(2) log(<i>Emails</i>)	(3) log(<i>Sales</i>)
2SLS:			
Log(Manager-Salary)	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)
Log(Peer-Salary)	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)
Reduced Form:			
(Log(Manager Signal)-Log(Manager Prior))*I(Manager Treatment)	0.103* (0.053)	0.094*** (0.029)	0.094 (0.066)
(Log(Peer Signal)-Log(Peer Prior))*I(Peer Treatment)	-0.379** (0.160)	-0.211* (0.112)	-0.381* (0.195)
First Stage Log(Manager-Salary):			
(Log(Manager Signal)-Log(Manager Prior))*I(Manager Treatment)	0.705*** (0.066)	0.692*** (0.035)	0.628*** (0.029)
(Log(Peer Signal)-Log(Peer Prior))*I(Peer Treatment)	0.018 (0.133)	0.073 (0.050)	0.125** (0.054)
First Stage Log(Peer-Salary):			
(Log(Manager Signal)-Log(Manager Prior))*I(Manager Treatment)	0.002 (0.011)	-0.008 (0.008)	-0.038*** (0.009)
(Log(Peer Signal)-Log(Peer Prior))*I(Peer Treatment)	0.404*** (0.115)	0.511*** (0.059)	0.538*** (0.043)
Observations	602	2,060	791

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. The panel *2SLS* reproduces the regression results from Table 2 – see its note for details about the variables and the specification. The remaining three panels reproduce the reduced form and first stage results from those Two-Stage Least Squares (2SLS) regressions. *Log(Manager-Salary)* and *Log(Peer-Salary)* correspond to the posterior beliefs about manager and peer average salaries, respectively. *Log(Manager Prior)* and *Log(Peer Prior)* correspond to the prior beliefs about manager and peer average salaries, respectively. *Log(Manager Signal)* and *Log(Peer Signal)* correspond to the signals about manager and peer average salaries, respectively. *I(Manager Treatment)* is a binary variable that takes the value 1 if the signal about manager salary was shown to the respondent and 0 otherwise. *I(Peer Treatment)* is a binary variable that takes the value 1 if the signal about peer salary was shown to the respondent and 0 otherwise.

Table C.12: Effects of Salary Perceptions on Effort and Performance: Alternative Specification

	Baseline Specification (logs)			Alternative Specification (%- Δ)		
	Hours (1)	Emails (2)	Sales (3)	Hours (4)	Emails (5)	Sales (6)
Post-Treatment (3-Months):						
<i>Log (Manager-Salary)⁽ⁱ⁾</i>	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)	0.523 (0.368)	0.465 (0.299)	0.552*** (0.210)
<i>Log (Peer-Salary)⁽ⁱⁱ⁾</i>	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	-3.210 (2.256)	-1.278** (0.591)	-1.992*** (0.395)
P-value H_0 : (i)=(ii)	0.026	0.007	0.000	0.141	0.016	0.000
Cragg-Donald F-Stat.	29.8	204.0	98.2	29.3	203.7	98.1
Observations	602	2,060	791	602	2,060	791

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different IV regression. Columns (1) through (3) are a reproduction of the baseline results from the top panel of Table 2 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to that table. The specifications in columns (4) through (6) are identical to columns (1) through (3) except that, instead of using the logarithm transformation, the dependent variable are defined as the ratio between the post-treatment outcome and the pre-treatment outcome.

Table C.13: Effects of Salary Perceptions on Effort and Performance: Additional Controls

	log(<i>Hours</i>) (1)	log(<i>Emails</i>) (2)	log(<i>Sales</i>) (3)	log(<i>Hours</i>) (4)	log(<i>Emails</i>) (5)	log(<i>Sales</i>) (6)	log(<i>Hours</i>) (7)	log(<i>Emails</i>) (8)	log(<i>Sales</i>) (9)	log(<i>Hours</i>) (10)	log(<i>Emails</i>) (11)	log(<i>Sales</i>) (12)
Post-Treatment (3-Months):												
<i>Log</i> (Manager-Salary) ⁽ⁱ⁾	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)	0.152** (0.075)	0.130*** (0.041)	0.106 (0.126)	0.153** (0.073)	0.119*** (0.040)	0.116 (0.137)	0.139 (0.089)	0.146*** (0.047)	0.067 (0.133)
<i>Log</i> (Peer-Salary) ⁽ⁱⁱ⁾	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	-1.010** (0.508)	-0.434** (0.210)	-0.720** (0.301)	-0.894* (0.536)	-0.414* (0.247)	-0.495* (0.286)	-0.329 (0.453)	-0.456** (0.213)	-0.872*** (0.289)
Gender Dummies				✓	✓	✓						
Location Dummies							✓	✓	✓			
Position Dummies										✓	✓	✓
P-value H_0 : (i)=(ii)	0.026	0.007	0.000	0.028	0.007	0.000	0.057	0.028	0.001	0.338	0.003	0.000
Cragg-Donald F-Stat.	29.8	204.0	98.2	29.6	204.0	98.3	24.4	198.2	94.5	41.7	253.2	92.8
Observations	602	2,060	791	602	2,060	791	602	2,060	791	602	2,060	791

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different IV regression. Columns (1) through (3) are a reproduction of the baseline results from the top panel of Table 2 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to that table. The specifications in columns (4)–(6) are identical to columns (1)–(3) except for the inclusion of an additional control variable: a binary variable indicating if the employee is female. Columns (7)–(9) are identical to columns (1)–(3) except for the inclusion of the following additional control variables: a set of 29 indicator variables corresponding to the location where the employee works. Columns (10)–(12) are identical to columns (1)–(3) except for the inclusion of the following additional control variables: a set of 180 indicator variables corresponding to the employee’s position title.

Table C.14: Effects of Salary Perceptions on Effort and Performance: Additional Controls

	$\log(\text{Hours})$	$\log(\text{Emails})$	$\log(\text{Sales})$	$\log(\text{Hours})$	$\log(\text{Emails})$	$\log(\text{Sales})$	$\log(\text{Hours})$	$\log(\text{Emails})$	$\log(\text{Sales})$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-Treatment (3-Months):									
<i>Log (Manager-Salary)⁽ⁱ⁾</i>	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)	0.152** (0.077)	0.129*** (0.040)	0.086 (0.095)	0.137* (0.074)	0.123*** (0.042)	0.128 (0.150)
<i>Log (Peer-Salary)⁽ⁱⁱ⁾</i>	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	-1.171** (0.496)	-0.499** (0.201)	-0.555* (0.306)	-1.105** (0.558)	-0.434** (0.214)	-0.729** (0.330)
Prior Gaps Dummies				✓	✓	✓			
Posterior Certainty							✓	✓	✓
P-value H_0 : (i)=(ii)	0.026	0.007	0.000	0.010	0.002	0.013	0.031	0.008	0.001
Cragg-Donald F-Stat.	29.8	204.0	98.2	31.0	196.0	95.4	26.7	198.7	97.8
Observations	602	2,060	791	602	2,060	791	597	2,035	782

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different IV regression. Columns (1) through (3) are a reproduction of the baseline results from the top panel of Table 2 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to that table. The specifications in columns (4)–(6) are identical to columns (1)–(3) except for the addition of the following control variables: sets of dummies indicating the decile of the prior gap in beliefs (nine dummies for manager salary and nine dummies for peer salary). Columns (7)–(9) are identical to columns (1)–(3) except for the addition of eight control variables (four variables for the manager salary and four variables for the peer salary). For the manager salary, the four variables are: the perceived probability that the true average manager salary falls between -10% and -2.5% of the posterior belief of the average manager salary; the perceived probability that the true salary falls between -2.5% and +2.5% of the posterior belief; the perceived probability that the true salary falls between +2.5% and +10% of the posterior belief; the perceived probability that the true salary falls above +10% of the posterior belief. And the four variables on peer salary are defined in an analogous way to the four variables on manager salary.

Table C.15: Effects of Perceived Manager and Peer Salary on Various Email Outcomes

	By Direction		Sent to		Sent to Pay Band		
	Sent (1)	Received (2)	Inside (3)	Outside (4)	Higher (5)	Same (6)	Lower (7)
<i>Log</i> (Manager-Salary)	0.130*** (0.041)	0.041 (0.035)	0.139*** (0.042)	0.047 (0.072)	0.249*** (0.086)	0.123* (0.065)	-0.047 (0.145)
<i>Log</i> (Peer-Salary)	-0.431** (0.210)	-0.198** (0.092)	-0.400* (0.207)	-0.602 (0.394)	0.046 (0.320)	-0.061 (0.276)	-1.076 (0.666)
P-value Diff.							
Manager		0.096		0.268		0.161	
Peer		0.310		0.650		0.464	
Observations	2,060	2,060	2,060	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 corresponds to a different IV regression. We always use the baseline specification from Table 2 – for more details about this specification see the notes to that table. Each column uses a different dependent variable: the average number of emails sent in the 90 days following the survey (column (1)), the number of emails received (column (2)), the number of emails sent to email accounts inside the firm (column (3)), emails sent outside the firm (column (4)), emails sent to employees with a higher (column (5)), same (column (6)) and lower (column (7)) paybands.

Table C.16: Heterogeneity in Effects of Perceptions on Number of Emails Sent

	(1)	(2)	(3)	(4)	(5)	(6)
	Female	> 4ys	High-Band	Sales	Front-Office	High Rank
$\text{Log}(\text{Manager-Salary})^{(i)}$	0.155*** (0.052)	0.075 (0.073)	0.129* (0.069)	0.196** (0.083)	0.153** (0.077)	0.118 (0.088)
$\text{Log}(\text{Peer-Salary})^{(ii)}$	-0.448** (0.224)	-0.215 (0.312)	-0.380 (0.277)	-0.548* (0.283)	-0.422* (0.231)	-0.321 (0.366)
Observations	1,506	1,054	898	972	1,454	750
	Male	≤ 4ys	Low-Band	Non-Sales	Back-Office	Low Rank
$\text{Log}(\text{Manager-Salary})^{(iii)}$	0.066 (0.092)	0.205*** (0.074)	0.127* (0.071)	0.130** (0.053)	0.094 (0.065)	0.149** (0.058)
$\text{Log}(\text{Peer-Salary})^{(iv)}$	-0.724 (0.707)	-0.616 (0.474)	-0.463 (0.342)	-0.135 (0.387)	-0.558 (0.613)	-0.443 (0.270)
Observations	554	1,006	1,162	1,088	606	1,310
P-value $H_0 : (i)=(iii)$	0.710	0.479	0.851	0.389	0.835	0.788
P-value $H_0 : (ii)=(iv)$	0.394	0.213	0.991	0.505	0.559	0.774

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different IV regression. We always use the baseline specification from Table 2 – for more details about this specification see the notes to that table. Each column corresponds to a different split of the sample (e.g., females vs. males), and the p-values correspond to the test of the null hypothesis that the learning rates are equal across the two subgroups. In column (6), there is missing data on perceived productivity rank for 52 observations, which we impute using the employee’s true position in the productivity distribution.

Table C.17: Asymmetry of Effects of Perceptions about Manager and Peer Salary: Upwards vs. Downwards Revisions

	Effort and Performance			Career Moves			
	log(<i>Hours</i>) (1)	log(<i>Emails</i>) (2)	log(<i>Sales</i>) (3)	P(Left) (4)	P(Transfer) (5)	log(<i>Salary</i>) (6)	$P(\Delta Title)$ (7)
Symmetric Model:							
<i>Log</i> (Manager-Salary)	0.150** (0.074)	0.130*** (0.041)	0.106 (0.122)	-0.016 (0.023)	-0.003 (0.030)	0.002 (0.011)	0.012 (0.029)
<i>Log</i> (Peer-Salary)	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	0.232** (0.106)	0.093 (0.106)	0.004 (0.052)	0.114 (0.123)
Asymmetric model:							
<i>Log</i> (Manager-Salary) Upwards ⁽ⁱ⁾	0.095 (0.078)	0.154*** (0.054)	0.141 (0.120)	-0.026 (0.025)	-0.025 (0.036)	-0.006 (0.014)	0.018 (0.029)
Downwards ⁽ⁱⁱ⁾	0.805 (0.542)	0.008 (0.255)	0.016 (0.227)	0.057 (0.075)	0.138 (0.101)	0.065 (0.048)	-0.042 (0.107)
<i>Log</i> (Peer-Salary) Upwards ⁽ⁱⁱⁱ⁾	0.051 (0.887)	-0.969** (0.480)	0.081 (0.257)	0.254* (0.154)	0.316 (0.245)	-0.064 (0.086)	0.194 (0.230)
Downwards ^(iv)	-1.930** (0.798)	0.094 (0.373)	-1.426*** (0.420)	0.228 (0.189)	-0.107 (0.131)	0.092 (0.061)	0.030 (0.081)
P-value Test: H_0 : (i)=(ii)	0.144	0.153	0.002	0.921	0.179	0.164	0.482
P-value Test: H_0 : (iii)=(iv)	0.204	0.612	0.621	0.313	0.152	0.186	0.580
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 corresponds to a different pair of IV regressions. In “Symmetric Model” we use the baseline specification from Table 2 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to that table. In “Asymmetric Model”, we use the specification described in Appendix C.15, which differs from the baseline specification in that it allows the effects of *Peer Salary* and *Manager Salary* to be different depending on whether the individual revised her prior beliefs upwards or downwards.

Table C.18: Asymmetry of Effects of Perceptions about Peer Salary: Above vs. Below Own Salary

	Effort and Performance			Career Moves			
	$\log(\text{Hours})$ (1)	$\log(\text{Emails})$ (2)	$\log(\text{Sales})$ (3)	P(Left) (4)	P(Transfer) (5)	$\log(\text{Salary})$ (6)	$P(\Delta\text{Title})$ (7)
Symmetric Model:							
<i>Log</i> (Peer-Salary)	-0.943** (0.472)	-0.431** (0.210)	-0.731** (0.297)	0.232** (0.106)	0.093 (0.106)	0.004 (0.052)	0.114 (0.123)
Asymmetric Model:							
<i>Log</i> (Peer-Salary)							
Above Own-Salary ⁽ⁱ⁾	-0.448 (0.600)	-0.402 (0.341)	-0.304 (0.276)	0.489*** (0.184)	-0.062 (0.162)	0.011 (0.077)	0.083 (0.140)
Below Own-Salary ⁽ⁱⁱ⁾	-2.357 (2.582)	-0.421 (0.392)	-1.114*** (0.388)	-0.099 (0.179)	0.291 (0.290)	0.023 (0.068)	0.151 (0.191)
P-value Test: $H_0: (i)=(ii)$	0.518	0.974	0.078	0.047	0.383	0.911	0.750
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 corresponds to a different pair of IV regressions. In “Symmetric Model” we use the baseline specification from Table 2 – for more details about the regression specification, such as outcomes, independent variables and data definitions, see the notes to that table. In “Asymmetric Model”, we use the specification described in Appendix C.15, which differs from the baseline specification in that it allows the effect of *Peer Salary* to be different depending on whether *Peer Salary* is below or above *Own-Salary*.

Table C.19: Effects of Salary Perceptions on Additional Survey Outcomes

	(1)	(2)
	Prod.	P(Prom.)
<i>Log</i> (Manager-Salary) ⁽ⁱ⁾	0.000 (0.015)	
<i>Log</i> (Peer-Salary) ⁽ⁱⁱ⁾	0.044 (0.040)	-0.014 (0.118)
P-Value (i)=(ii)	0.280	
Cragg-Donald F-Stat.	250.5	512.1
Mean Dep. Var.	0.47	0.56
Std. Dev. Dep. Var.	0.22	0.29
Observations	1,999	2,051

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for a different IV regressions, following the same specification from Table 3 – see the corresponding notes for more details. *Manager-Salary* is the posterior belief about manager salary, and *Peer-Salary* is the posterior belief about the average peer salary. *Prod.* denotes the individual self-perceived position in the distribution of performance ratings in the firm and was elicited after the elicitation of the posterior beliefs. *P(Prom.)* is the perceived probability of being promoted to the managerial position in the next five years. Since this belief was elicited prior to the experimental provision of information about manager salary, we do not include *Manager-Salary* as independent variable.

Table C.20: Determinants of Horizontal Variation in Salaries

	Dep. Var.: $\log(\text{Salary}) - \log(\text{Average Peer Salary})$					
	(1)	(2)	(3)	(4)	(5)	(6)
Within-Peer-Group Rank in:						
Performance Rating	0.052*** (0.007)					0.035*** (0.010)
Hours Worked		0.016 (0.013)				
Emails Sent			0.055** (0.022)			
Sales				0.064** (0.023)		
Demographics:						
Female					-0.013** (0.006)	-0.013** (0.006)
Age					0.005*** (0.001)	0.005*** (0.001)
$\log(\text{Tenure})$					0.024*** (0.007)	0.023*** (0.007)
College					-0.001 (0.007)	-0.003 (0.007)
Business Major					0.002 (0.008)	0.003 (0.008)
Finance Major					0.008 (0.007)	0.009 (0.006)
Mean Outcome	-0.011	-0.009	-0.013	-0.010	-0.011	-0.011
Std. Dev. Outcome	0.157	0.141	0.159	0.144	0.157	0.157
R^2	0.009	0.001	0.011	0.017	0.073	0.076
Observations	(Masked)	1,765	4,587	1,597	(Masked)	(Masked)

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for a separate OLS regression. The dependent variable is the log-difference between the individual's own base salary and the average salary in the peer group. *Performance Rating*, *Hours Worked*, *Emails Sent* and *Sales* correspond to the individual's rankings in their peer groups and thus range from 0 (bottom of the group) to 1 (top of the group). *Performance Rating* corresponds to the ranking of performance review rating in 2016. *Hours worked* correspond to the ranking in the number of hours worked in the three months prior to the experiment. *Emails Sent* correspond to the ranking in the number of emails sent in the three months prior to the experiment. *Sales* correspond to the ranking in the sales performance index during the three months prior to the experiment. *Female* is an indicator variable that takes the value 1 if the employee is female. *Age* is the age in years. $\log(\text{Tenure})$ is the (log) of the number of months since the employee joined the firm. *College* takes the value 1 if the employee has a College degree. *Business Major* takes the value 1 if the employee's College degree is a business major. *Finance Major* takes the value 1 if the employee's College degree is a finance major. Columns (1), (5) and (6) includes all employees working at the firm at the start of the experiment. Columns (2) through (4) correspond to the subsample of employees with non-missing data in *Hours Worked*, *Emails Sent* and *Sales*, respectively.

Table C.21: Average Effects of Information Disclosure

	Effort and Performance		
	(1)	(2)	(3)
	log(<i>Hours</i>)	log(<i>Emails</i>)	log(<i>Sales</i>)
Post-Treatment:			
Shown Manager Info ⁽ⁱ⁾	0.040 (0.039)	0.033 (0.026)	0.026 (0.032)
Shown Peer Info ⁽ⁱⁱ⁾	0.044 (0.056)	-0.020 (0.036)	-0.009 (0.021)
Pre-Treatment (Falsification):			
Shown Manager Info	-0.012 (0.043)	0.014 (0.024)	0.032 (0.039)
Shown Peer Info	0.010 (0.045)	-0.006 (0.030)	-0.029 (0.032)
P-value H_0 : (i)=(ii)	0.959	0.231	0.426
Mean Outcome	5.98	35.57	0.48
Std. Dev. Outcome	1.88	44.93	0.23
Observations	602	2,060	791

Notes: Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column presents results for two sets of OLS regressions, each based on equation (C.7) from Appendix C.19: 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. *Shown Manager Info* is a binary variable indicating whether the subject was randomly chosen to receive information about the average manager salary. *Shown Peer Info* is a binary variable indicating whether the subject was randomly chosen to receive information about the average peer salary. The regressions control for three monthly lags of the dependent variable, (log) own salary, (log) tenure, 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. The mean and std. dev. reported in the bottom rows correspond to the values prior to taking the logarithm function. Columns (1) corresponds to the subsample of employees in the headquarter offices and column (3) to the subsample of employees with sales roles.