

How Much Does Your Boss Make? The Effects of Salary Comparisons

Zoë Cullen

Harvard Business School

Ricardo Perez-Truglia

University of California, Berkeley

The vast majority of the pay inequality in organizations 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 document large misperceptions among employees about the salaries of their managers and smaller but still significant misperceptions of the salaries of their peers, and we show that these perceptions have a significant causal effect on the employees' own behavior.

I. Introduction

The vast majority of 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

We thank the editor and the anonymous referees for their very useful feedback. We are thankful for excellent comments from numerous colleagues and seminar participants. The collaborating institution provided financial support for the research being conducted. Additionally, Zoë 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 the University of California, Los Angeles (IRB 17-001529). This paper was edited by John A. List.

Electronically published January 19, 2022

Journal of Political Economy, volume 130, number 3, March 2022.

© 2022 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/717891>

unclear, however, whether employees are aware of pay inequality in their own firms. Pay secrecy policies, for example, may prevent employees from discovering how much their managers are making. Even if employees knew that their managers earned substantially more, precisely how that knowledge would affect employees' behavior remains unknown. On the one hand, learning that managers earn a lot may demoralize employees as a result of social preferences; they may feel jealous or resentful, or they may perhaps feel that the pay disparity is unfair. On the other hand, employees may find well-paid managers to be a source of motivation, an extra incentive to work harder in the hope 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 that their managers are paid handsomely, does that inspire them to put in a 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 elicitation are incentivized to ensure that giving thoughtful, honest answers is in the best interest of participants. To measure the respondents' misperceptions about salaries, we can compare their perceptions with the true average salaries listed in 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. To assess whether they even care about others' salaries, we measure their willingness to pay for information, using an information-acquisition experiment (Becker, DeGroot, and Marschak 1964).

Central to our research design is the information-provision experiment. After eliciting prior beliefs, we randomized whether the employee would receive a signal about salary. For example, after eliciting beliefs about 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 reelicited beliefs regardless of whether the subject received the signal. 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 employees' behavior as measured by administrative records. This method allows us to test whether higher perceived manager salary translated into higher effort, consistent with career concerns, or demotivated employees, consistent with social preferences at play.

Moreover, our research design allows us to provide some suggestive evidence of 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 toward pay inequality. Our research design allows for an additional test of career concerns. According to our model of career concerns, the effects of manager salary should depend on whether the employee can aspire to be promoted to that position. 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). To complement these data, we elicited the respondent's own perceptions of the number of promotions they would need to secure the managerial position and the likelihood that they will be promoted to that position in the next 5 years.

We conducted the natural field experiment (Harrison and List 2004) with a sample of 2,060 employees from a large commercial bank (hereafter referred to as "the firm") located in a lower-middle-income country in Southeast Asia. The firm is comparable to other firms around the world in two key respects: pay inequality and lack of pay transparency. The firm discloses limited information about pay, and according to our data, employees at this firm rarely discuss salaries with their coworkers. Abundant evidence indicates that both circumstances are the reality in most firms across the world, including but not limited to the United States (Trachtman 1999; Edwards 2005; Hegewisch, Williams, and Drago 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 the 90th percentile of salaries is 0.21 in the firm compared with 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, Jensen, and Murphy 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 have significant misperceptions about the salaries of their managers and provide some evidence of the sources of this information friction. We find that employees' guesses about the average manager salary have a mean absolute error of 28%. Moreover, a systematic bias is at work: 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 that employees do not have information in the first place: lack of social learning. We rolled out the survey gradually over the course of 2 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 both underestimate the average manager salary by 20% at the start of the survey. 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—say, by 10%. 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 to 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 ($p = .042$). The corresponding effects on the other measures of effort and performance are similar in magnitude: elasticities of 0.130 ($p = .001$) in the number of emails sent and 0.106 ($p = .383$) in sales performance. The effects of the horizontal comparison instead go in the opposite direction to 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 ($p = .045$), with corresponding elasticities of -0.431 ($p = .041$) in emails sent and -0.731 ($p = .014$) in sales performance. Indeed, we can confidently reject the null hypothesis that the effects of horizontal and vertical comparisons are equal to each other: $p < .026$ for hours worked, $p = .007$ for emails sent, and $p < .001$ 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 career concerns, which predict positive effects, dominate over social preferences, which predict 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, consistent with career concerns, when employees learn that their managers earn more, they become more optimistic about what their own salaries will be 5 years in the future. Alternatively, and contrary to what our model of social preferences channel would predict, we do not find any evidence that perceptions of manager salary have an 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 has 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 career concerns, 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 5 years and they work harder. In contrast, when employees discover 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 for 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 for which we provide evidence indicate that the results should change based on some mediating factors. For example, take 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 in which 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 up 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 et al. (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 increased the intention to switch jobs. More recently, other studies have documented effects of pay inequality and pay transparency using natural experiments (Mas 2017; Dube, Giuliano, and Leonard 2019; Perez-Truglia 2020), field experiments (Cohn et al. 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 literature in two ways. First, whereas the previous work focuses 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, Jensen, and Murphy 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. Likely 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 US companies. He finds that 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 (CEOs) 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 the literature about career concerns and salary dynamics (Lazear and Rosen 1981; Harris and Holmstrom 1982; Rosen 1986; Gibbons and Murphy 1992; Dewatripont, Jewitt, and Tirole 1999; Gibbons and Waldman 1999a, 1999b; Holmstrom 1999). Although

a rich body of theoretical work looks at 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 on the basis of 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 et al. 2019) and on the effects of relative income on job satisfaction (Clark, Frijters, and Shields 2008; Godechot and Senik 2015) and happiness (Senik 2004).

The rest of the paper proceeds as follows. Section II describes the research design, including the institutional context and data sources. Section III presents the results on salary misperceptions. Section IV discusses the effects of perceived salaries on behavior. Section V discusses the generalizability of the results and its implications for pay transparency. Section VI concludes.

II. Research Design

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

A. Conceptual Framework

Why would employees want to be informed about the salaries of their managers and peers? How would that information affect their own 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 2-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: $c(e) = (1/2)e^2$.

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: $\tilde{w}^z(\bar{w}^{\text{mgr}})$, with $\partial \tilde{w}^z / \partial \bar{w}^{\text{mgr}} > 0$.¹ 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

¹ We assume that 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 ($\tilde{w}^r > \tilde{w}^z$).

average salary of peers. We assume that the expected raise if not promoted is an increasing function of the average salary of peers: $\tilde{w}^r(\bar{w}^{\text{peer}})$, with $\partial\tilde{w}^r/\partial\bar{w}^{\text{peer}} > 0$. 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.

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

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

Employees choose effort to maximize expected utility given by equation (1). Our functional form assumptions ensure that the objective function is concave, $U_{ee} < 0$, and thus that 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^*):

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

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:

$$\frac{\partial e^*}{\partial w^{\text{mgr}}} = \underbrace{\delta \cdot \theta \cdot \frac{\partial \tilde{w}^z}{\partial w^{\text{mgr}}}}_{\text{career concerns}} \underbrace{- \psi^{\text{mgr}}}_{\text{social preferences}}. \quad (3)$$

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:

² There 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.

$$\frac{\partial e^*}{\partial w^{\text{peer}}} = \underbrace{-\delta \cdot \theta \cdot \frac{\partial \tilde{w}^r}{\partial w^{\text{peer}}}}_{\text{career concerns}} \underbrace{-\psi^{\text{peer}}}_{\text{social preferences}}. \quad (4)$$

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 refer to career concerns and social preferences as two distinct channels. On the one hand, keep in mind that this is a toy model and thus not intended to capture all of 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 of or against these specific mechanisms. First, the career concerns channel predicts that a higher manager salary should increase one's own expected 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 they are learning about.

B. 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 maintain the confidentiality of the firm's identity, we refrain from being specific about its characteristics. This firm has millions of customers, billions of dollars in assets and revenues, and thousands of employees. These employees are based in two headquarters and in hundreds of branches dispersed around 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 the 90th percentile of salaries is 0.21 in this firm, whereas it is 0.19 for the average medium-sized US firm (Song et al. 2019). The inequality in this firm is also typical in that only a small part of it is horizontal. A

³ In 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 that their peers are paid better. Indeed, Caldwell and Harmon (2018) provide some evidence in support of this mechanism.

⁴ There are two reasons for that result. The direct effect is that salary expectations are conditional on a 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.

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 (Baker, Gibbs, and Holmstrom 1994; Card et al. 2012).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 US employees at Bank of America.

Regarding pay transparency, the firm does not have an open salary policy.⁶ Several 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 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). For example, in a survey of private-sector employees from the United States, more than 60% report that their employer discourages or prohibits employees from discussing salaries with coworkers (Hegewisch, Williams, and Drago 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 shows that 70% report never discussing salaries with coworkers (Gee 2017).

Among the observable characteristics, some features of this firm 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 as well. We return to this topic in section V.A, where we discuss how the results from this specific firm should be extrapolated to other settings.

⁵ For details, see app. C.1 (apps. A–C are available online).

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

⁷ More precisely, 45% of employees reported never talking about salaries, 16% reported talking once a year, 31% reported talking a few times a year, 6% reported talking once a month, and the remaining 2% reported talking once a week or more often. Because the employer frowns on such discussions, they are probably underreported in surveys.

⁸ At the time of the experiment, the share of female employees at this firm was 71%. In comparison, at around the same time that we conducted our experiment, the US workforce of Bank of America was 53% female (see <https://about.bankofamerica.com/en-us/what-guides-us/our-global-workforce.html>).

C. *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 whom 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, and (3) the managerial position had to have an oversight role over the respondent, such as conducting performance evaluations or approving leaves of absence.

According to the career concerns mechanism from section II.A, 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 larger number of promotions away). To complement these 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 5 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; Cullen and Pakzad-Hurson 2016; Breza, Kaur, and Shamdasani 2018; Dube, Giuliano, and Leonard 2019): 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.

D. *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

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

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 this paper have been converted to US dollars using purchasing power parity (PPP)–adjusted exchange rates from April 2017.¹⁰ To ensure that subjects understood the incentives for truth-telling, in the training portion of the survey we explained that it is in 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 preclude 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

¹⁰ More 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 explicitly disclose the formula for accuracy rewards—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, Vesterlund, and Wilson 2020).

¹¹ We did not disclose the exact formula used for the accuracy rewards, thus 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 the total amount earned was 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. The average total reward for survey participation was around \$30, roughly half a day's salary for the typical subject.

¹² Using 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

the definition of salary, we included an exercise in which 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 to avoid revealing them to the researchers.

E. 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 assess 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 dispense this information in an incentive-compatible way, we employed the multiple price list variation of the Becker-DeGroot-Marschak (BDM) method (Becker, DeGroot, and Marschak 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). Because 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: \$1.30, \$6.50, \$26.10, \$130.50, and \$652.30.¹⁴ 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 implemented. For the 1% of respondents who had their choice implemented (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.

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 we excluded those employees from participation in our study.

¹³ We 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 reexplaining the definition.

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

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 free. In the following screen, we let the subjects know whether they were chosen to receive the signal.

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 relicit 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, followed by the module for manager salary. With respect to the information-provision experiment, we cross-randomized the two pieces of information, generating 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 their 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 in which employees are most interested. 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

¹⁵ When 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.

F. 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 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 (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 access their work email account only from their office computers, implying that they can send emails only 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, owing 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, a recent but growing literature in information science demonstrates that the email behavior predicts employee performance (for a literature review, see Wen et al. 2020).

¹⁶ Since some employees work on weekends, we average over all 7 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 US Bureau of Labor Statistics, the average American works 44 hours per week (CNBC 2017). Divided by 7 days of the week, the daily average is 6.29 hours, the same order of magnitude as the corresponding average in our sample (5.98 hours, from table 2, col. 1).

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

¹⁸ Results are reported in app. C.2.

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, 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 that our measures of effort—hours in the office and emails sent—are likely to be elastic as well. Employees seem to have quite a bit of discretion with how long to work (leaving a bit earlier or staying after hours) and with how many emails to send and respond to. The company does not explicitly monitor the hours worked and emails sent, nor does it reward or punish employees 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 zero (lowest) to one (highest). As a validation exercise, we show that the number of emails is positively and significantly associated with 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 the human resources department to track those secondary outcomes. We began collecting all the administrative data 3 months before launching the survey. As a result, in addition to posttreatment outcomes, we can measure the pretreatment outcomes, which can be used as control variables to improve precision as well as for falsification tests.

G. 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 the 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 II.A, 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

¹⁹ Results are reported in app. C.2.

about respondents' own future salaries 1 and 5 years in the future, using an incentive-compatible method.²⁰

According to the social preferences channel from section II.A, the salary information may affect employee morale and tolerance for pay inequality. We included three questions meant to gauge this mechanism. Following previous 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?" Responses to this question are on a five-point scale from very dissatisfied (1) to very satisfied (5). The second question, about overall job satisfaction, uses the same five-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, Perez-Truglia, and Tetaz 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.

H. 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 reward 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

²⁰ To incentivize truthful responses, offering rewards by comparing the guesses with the actual future salaries was not practical because we would need to wait 5 years to calculate the rewards. Instead, we told respondents that we would compare their guesses with our own predictions of their future salaries (based on our administrative data and predictive models).

²¹ The 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 because of 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.

²² The 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.

were sent out gradually over the course of 2 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

I. 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 5 years. Seventy-three percent 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 provides 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.

Some additional features of the setting are particularly important to keep in mind. First, while the salary differences between peers tends to be small, large differences exist 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, the company has plenty of opportunities for upward mobility (the annual promotion rate is around 16.5%), of which employees seem to be well aware.

III. 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.

A. 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 1A shows misperceptions of average manager salary. Only a small share

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

²⁴ This 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).

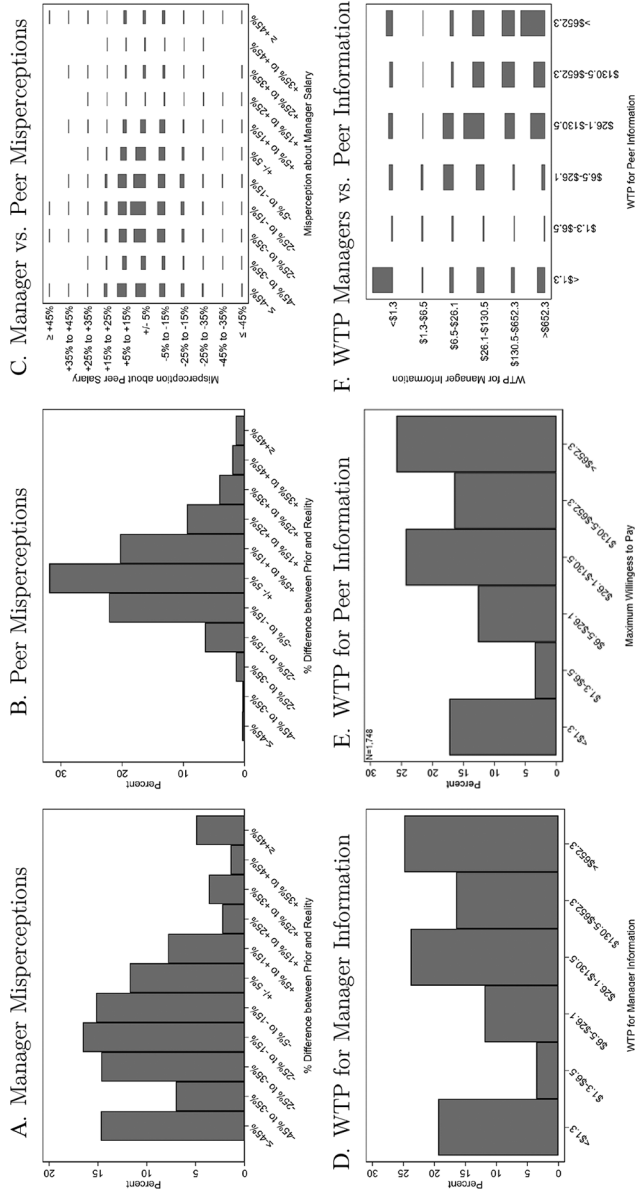


FIG. 1.—Salary misperceptions and willingness to pay (WTP) for salary information. *A–C* correspond to salary misperceptions defined as the employee’s prior belief (according to an incentivized survey question) and the actual salary (according to the firm’s administrative records) divided by the actual salary ($N = 2,060$). *A* focuses on the average manager salary, *B* focuses on the average peer salary, and *C* focuses on the joint distribution of *A* and *B*. *D* and *E* show the WTP for a specific information piece (an average based on a sample of five managers’/peers’ salaries) based on the responses to multiple price list questions. *D* corresponds to the WTP for manager information ($N = 1,637$ respondents with consistent responses across the five scenarios). *E* corresponds to the WTP for peer information ($N = 1,748$ respondents with consistent responses across the five scenarios). *F* corresponds to the joint distribution of *D* and *E* ($N = 1,478$ respondents with consistent responses for both the manager and the peer scenarios).

(12%) of respondents guess the average manager salary within $\pm 5\%$ accuracy. The rest of the respondents miss the mark, often by a large margin: the mean absolute error is 28%. As a benchmark, figure 1B 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 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, though of only 2.5% ($p < .01$).

We can also provide some robustness checks. 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 1C, the misperceptions of manager and peer salary are largely unrelated to each other: their correlation coefficient is statistically significant ($p = .007$) 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 have precise knowledge of the manager and peer salaries. For example, respondents on average believe that there is a 32.2% probability that the true manager salary falls within 2.5% (i.e., $\pm 2.5\%$) 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 finding as evidence that, despite being aware that they are far from perfectly accurate, employees are overconfident about their accuracy.

²⁵ Results are reported in app. C.4.

²⁶ Cullen 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 vs. 14.6% in Cullen and Perez-Truglia (2018).

²⁷ Results are reported in app. C.5.

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; moreover, those who do not report their own salary are not faring any better than if they had reported their own salary.²⁸ While reporting one's own salary is a reasonable idea when guessing a 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.

B. Willingness to Pay for Salary Information

The misperceptions of the salaries of peers and managers may simply reflect a lack of interest in 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 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 1D 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 very much. At one extreme, 22.9% of employees are not willing to pay more than \$6.50 for the signal about manager salary, an amount that is typically less than an hour's worth of salary.³⁰ On the other

²⁸ We 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).

²⁹ Responses 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, they 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; Fuster et al. 2018; Allcott and Kessler 2019).

³⁰ We 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 \$1.30. In principle, some of these subjects could have a negative willingness to pay—i.e., 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.

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 1E 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, substantial discrepancies arise at the individual level: as shown in figure 1F, plenty of employees are interested in the manager information but not the peer information and vice versa.³¹

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 with 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, Yim, and Prokopy 2003), \$0.80 for food certification information (Angulo, Gil, and Tamburo 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 nonincentivized 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)

³¹ More precisely, the correlation coefficient between the willingness to pay for manager and peer information is 0.28 ($p < .001$).

³² One 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 formula indicates an average willingness to pay of \$333 for the manager information and \$343 for the peer information.

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

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 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 II.A, the employee may need the information on the manager pay to decide whether to work harder to get promoted or to use it as a bargaining chip in future salary negotiations. Indeed, appendix C.7 provides some suggestive evidence in support of this interpretation.

C. 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 IV.

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, Cruces, and Perez-Truglia 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} represent 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 a value of one if we showed that signal to employee i and zero if not. Denote M_i^{post} as

³⁴ Results are reported in app. C.6.

³⁵ Theoretically, any positive probability of executing a choice ensures that the elicitation is incentive compatible; Carson and Groves (2007) and Charness, Gneezy, and Halladay (2016) provide empirical evidence in support of this conclusion. 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.

³⁶ In 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 (Fuster et al. 2018; Allcott and Kessler 2019).

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 zero (individuals ignore the signal) to one (individuals fully adjust to the signal), and according to the Bayesian model, α should depend on the relative precision between the prior belief and the signal. We must note that 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:

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

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 overestimated salaries would revise their beliefs downward when shown the signal, while those who underestimated salaries would revise their beliefs upward when shown the signal.

The learning results for manager salary are presented in the top half of figure 2. Figures 2A and 2B show 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 on manager salary, and the y -axis shows the revision observed in practice. Each panel includes a scatterplot of the raw data, with one dot per subject as well as a binned scatterplot and a linear fit.

Figure 2A corresponds to subjects who received information about the manager salary. If these subjects found the feedback useful, we should expect them to update their beliefs toward the feedback. As expected, figure 2A shows a strong relationship between the belief revisions and prior gaps: an additional percentage point (pp) in perception gap is associated with a revision that is 0.78 pp higher. There is some heterogeneity in the reaction to the feedback, as well: most employees lie close to the 45-degree line, implying that 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 a third group of employees lie in between the previous two groups, meaning that they adjusted their

³⁷ To 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 as if they are updating by the largest extent.

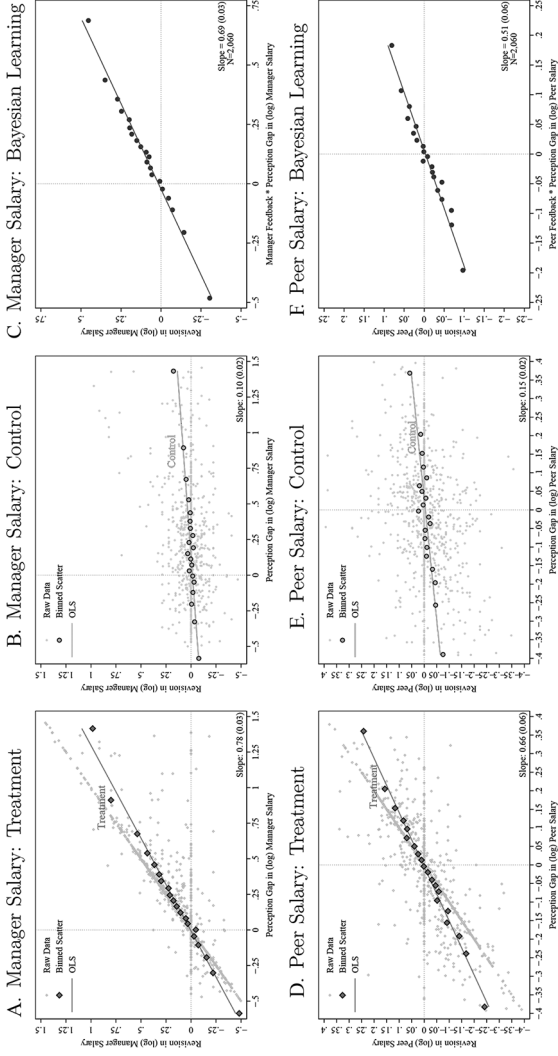


FIG. 2.—Effects of information-provision experiment on salary perceptions. $N = 2,060$. In *A* and *B*, the y -axis corresponds to the respondent’s update (i.e., the posterior belief about the average manager salary minus the corresponding prior belief) and the x -axis corresponds to the difference between the feedback chosen for the employee (the average salary among the random sample of five managers) and the respondent’s prior belief. *A* corresponds to subjects who were shown feedback about the manager salary. *B* is identical to *A*, except that it corresponds to subjects who were not selected to receive feedback about the manager salary. The raw data correspond to a regular scatterplot, where each circle/diamond corresponds to a different respondent (for ease of exposition, we do not plot data points outside the range $[-0.5, 1.5]$). The larger circles/diamonds labeled “binned scatter” correspond to a binned scatterplot based on the same data. The slopes correspond to a linear regression, with standard errors clustered at the position level and presented in parentheses. *C* estimates the Bayesian learning equation (6) from section III.C—this regression controls for the difference between the feedback chosen for the employee and the employee’s prior belief, and position title dummies. *D–F* are equivalent to *A–C* except that they are about peer salary instead of manager salary. OLS = ordinary least squares. A color version of this figure is available online.

posterior beliefs only partially. This heterogeneity may reflect differences in learning rates, survey attention, or simple measurement error.³⁸

Figure 2*B* corresponds to the subjects who did not receive information about the manager salary. Since these subjects did not get to see the feedback, we should not expect them to update toward the feedback. However, 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. For example, respondents may 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. As expected, figure 2*B* shows that the relationship between the belief updating and prior gaps 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.³⁹ This finding suggests a statistically significant ($p < .001$) but economically small degree of spurious revision. Indeed, this result is consistent in terms of magnitude with other information-provision studies (e.g., Cavallo, Cruces, and Perez-Truglia 2017; Fuster et al. 2018).

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

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

The parameter β picks up the spurious reversion toward the signal, while α picks up the true learning—that is, the degree of revision caused by the information provision, above and beyond the spurious revision. 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

³⁸ The Bayesian learning model predicts that different individuals may react differently to the signal depending on how confident they feel about it and their own prior beliefs. In the extreme case, if the employee was 100% sure about their own prior belief, Bayesian learning would suggest that such an 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. Recall that the signal is based on a random sample of five employees, which contains some sampling variation by construction. Additionally, some of the lack of updating may be due to inattention—e.g., owing 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.

³⁹ Some observations are somewhat above and somewhat below the x-axis, which means that some subjects revised their beliefs even when they did not receive any feedback. This outcome could be due to a number of reasons but most likely reflects the typical measurement error observed with survey measures of beliefs.

significantly larger than the precision of prior beliefs, we should expect α to be substantially above zero.⁴⁰

Figure 2C estimates equation (6), which yields the learning rate. The y -axis is still the revision 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 IV.A. Figure 2C 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 employee puts a weight of 69% on the signal provided by the experimenter and the remaining 31% on the prior belief.

Let P_i^{prior} , P_i^{signal} , P_i^{post} , and T_i^p represent, 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. The bottom half of figure 2 is identical to the top half of figure 2, 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. Individuals also learned significantly from the peer salary information: the slope from figure 2F 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.⁴¹

Appendix C.8 provides some additional results; it shows that learning from the feedback was compartmentalized (i.e., subjects did not use the information about peer salary to update beliefs about the manager salary), and it shows that learning rates were similar across different subsets of the population, such as between female and male employees.

D. 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 afford us insights about the sources of the misperceptions documented above. Additionally, the information diffusion is relevant for the

⁴⁰ For example, for manager salary, the mean absolute error of our signal is 6.8%, while the corresponding figure for prior beliefs is 28%.

⁴¹ While 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 ($p < .001$). According to the Bayesian model, this finding would indicate that individuals thought that the manager signal was more precise and/or felt more confident about their prior beliefs about peer salary.

interpretation of the experiment results: 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 formulate 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 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: $M_i^{\text{abs}} = |(M_i^{\text{post}} - M_i^{\text{true}})/M_i^{\text{true}}|$.⁴² The regression of interest is as follows:

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

The binary variable T_i^M indicates 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. The variable I_i^M takes a value of one 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 they started the survey).⁴³ For instance, in the baseline specification, I_i^M is a binary variable that takes a value of one if i 's closest peer received the information.⁴⁴ 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.⁴⁵

⁴² In app. C.9, we show that 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.

⁴³ In the baseline specification, we always define I_i^M to take a value of zero 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 app. C.9, we show that the results are robust under an alternative specification.

⁴⁴ The 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.

⁴⁵ Appendix C.9 provides descriptive statistics for all of the main variables used for the analysis of information diffusion.

The regression results are presented in table 1. Columns 1–4 correspond to misperceptions of manager salary.⁴⁶ 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, at 16 pp ($p < .001$). For reference, the average of the dependent variable in the control group is 27.6 pp, so this amounts to a 58% reduction in misperceptions.

Column 2 includes an additional variable related to indirect information provision: *Closest Peer* is a binary variable that takes a value of one 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 largest total of emails (sent and received) over the 3 months preceding the start of the experiment. Even though this measure is based on email data, it is plausible that it is also correlated with face-to-face interactions. For instance, data on card swipes confirm that these individuals go to lunch together more often than with other peers.⁴⁷

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*. However, the results indicate 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 ($p < .001$). 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

⁴⁶ One 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.

⁴⁷ For employees working in the headquarters offices, we can use the swipe data to proxy whether a given pair of employees have lunch together—i.e., 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).

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	-.160*** (.011)	-.159*** (.011)	-.168*** (.013)	-.168*** (.017)	-.156*** (.014)	-.045*** (.005)	-.044*** (.005)	-.042*** (.007)	-.043*** (.009)	-.044*** (.007)
Indirect: Closest Peer		.001 (.019)				.009 (.011)				
Indirect: Number of Peers			-.003 (.003)					.001 (.001)		
Indirect: Number of Peers >0				-.011 (.019)					.002 (.009)	
Indirect: Share of Peers					.031 (.069)					.012 (.028)
Constant	.276*** (.019)	.276*** (.019)	.276*** (.019)	.280*** (.021)	.274*** (.019)	.111*** (.011)	.111*** (.011)	.111*** (.011)	.110*** (.012)	.110*** (.011)
Observations	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060	2,060

NOTE.— $N = 2,060$. Standard errors (in parentheses) are clustered at the peer group level. Each column corresponds to a different regression. All regressions follow the econometric model described in sec. III.D. In cols. 1–5, the dependent variable is the absolute value of the difference between the posterior belief about average manager salary and the true average. “Direct” is a binary variable indicating whether the subject received the signal on manager salary. The other independent variables measure whether the peers of the subject received the information before the subject completed the survey. Those variables take a value of zero if the employee received the information directly. “Closest Peer” is a binary variable indicating whether the individual’s closest peer (defined as the peer with whom the employee exchanges the most number of emails in January–March 2017) received the information. “Number of Peers” is the number of peers who received the information. “Number of Peers >0” is a binary variable indicating whether at least one peer received the information. “Share of Peers” is the share of peers who received the 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 cols. 1–5 but using peer salary instead of manager salary.

*** Significant at 1%.

refrain from discussing salaries because of social norms around privacy (i.e., the “salary taboo”).⁴⁸

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. Number of Peers measures the number of peers who received information (before the respondent completed the survey, as always). Number of Peers >0 is a binary variable that takes a value of one if at least one peer received the information. 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 nonexperimental 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.

IV. Results: The Effects of Salary Beliefs on Behavior

A. Econometric Model

Let Y_i^{post} represent 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 III.C that M_i^{post} denotes the posterior belief about the average manager salary and P_i^{post} represents the posterior belief about the average peer salary. The following equation establishes the relationship of interest:

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

The parameter η_{mgr} captures the effects of manager salary. According to the model from section II.A, 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 ($\eta_{\text{mgr}} > 0$), but the higher manager salary could also demoralize the employee ($\eta_{\text{mgr}} < 0$). In turn, the parameter η_{peer} captures the effects of peer salary. According to the model

⁴⁸ Indeed, in a follow-up study (Cullen and Perez-Truglia 2018), we provide evidence in support of this channel.

from section II.A, both channels predict a negative effect ($\eta_{\text{peer}} < 0$): a higher peer salary mitigates the incentive to work hard for a promotion and at the same time may demoralize the employee because of 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 $\eta_{\text{mgr}} > 0$. 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 us 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 outcome 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 greater effort, that shift 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 these 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: $\eta_{\text{mgr}} = 0.2 = 2\%/10\%$.

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 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. The following IV regression is designed to do just that:

$$\begin{aligned}
\log(Y_i^{\text{post}}) &= \pi_0 + \eta_{\text{mgr}} \cdot \log(M_i^{\text{post}}) + \eta_{\text{peer}} \cdot \log(P_i^{\text{post}}) \\
&\quad + \pi_1 (\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \\
&\quad + \pi_2 (\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) \\
&\quad + \pi_3 \cdot \log(M_i^{\text{prior}}) + \pi_4 \cdot \log(P_i^{\text{prior}}) + X_i \pi_5 + \epsilon_i,
\end{aligned} \tag{9}$$

$$\begin{aligned}
\log(M_i^{\text{post}}) &= \mu_0 + \mu_1 (\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \cdot T_i^M \\
&\quad + \mu_2 (\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) \cdot T_i^P \\
&\quad + \mu_3 (\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \\
&\quad + \mu_4 (\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) + \mu_5 \cdot \log(M_i^{\text{prior}}) \\
&\quad + \mu_6 \cdot \log(P_i^{\text{prior}}) + X_i \mu_7 + \xi_i^1,
\end{aligned} \tag{10}$$

$$\begin{aligned}
\log(P_i^{\text{post}}) &= \nu_0 + \nu_1 (\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \cdot T_i^M \\
&\quad + \nu_2 (\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) \cdot T_i^P \\
&\quad + \nu_3 (\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \\
&\quad + \nu_4 (\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) + \nu_5 \cdot \log(M_i^{\text{prior}}) \\
&\quad + \nu_6 \cdot \log(P_i^{\text{prior}}) + X_i \nu_7 + \xi_i^2.
\end{aligned} \tag{11}$$

Equations (10) and (11) correspond to the first stage of the IV regression. These regressions measure the effect of the information-provision experiment on the posterior beliefs and are based on the Bayesian learning model given by equation (6) from section III.C. 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) below.

To determine the source of the identification, 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 estimating the same slope but for individuals in the treatment group (i.e., individuals who were randomly assigned to receive information). Slopes that are the same in the two groups would indicate that the salary information did not affect the outcome.

Slopes that are statistically different between the treatment and control groups 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.

Differences in slopes can be attributed to random assignment, but that does not prove that the effects of information can be attributed to the differences in posterior beliefs (i.e., learning), which is the key assumption made by the IV specification. In mathematical terms, the instrument exogeneity assumptions are $E[(\log(M_i^{\text{signal}}) - \log(M_i^{\text{prior}})) \cdot T_i^M \cdot \epsilon_i] = 0$ and $E[(\log(P_i^{\text{signal}}) - \log(P_i^{\text{prior}})) \cdot T_i^p \cdot \epsilon_i] = 0$. 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 described above, except 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 potential bias 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 toward zero.

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 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.

B. *Main Results*

Table 2 presents the results from the IV regression model.⁴⁹ 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 panel A, 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 1–3.⁵⁰ 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 $\log(\text{Manager Salary})$ is positive (0.150) and significant, both statistically ($p = .042$) and economically. This coefficient indicates that believing that their managers' salaries are higher, on average, motivates employees. Since the right- and left-hand-side variables are defined in logs, this coefficient implies a behavioral elasticity of 0.150; that is, 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.⁵¹ The coefficient on $\log(\text{Manager Salary})$ is positive (0.130) and highly significant ($p = .001$).⁵² 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; $p = .816$).⁵³ 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 $\log(\text{Manager Salary})$ from column 3 is positive (0.106) and on the same order of magnitude as the corresponding coefficients for the

⁴⁹ For the sake of brevity, this table presents the IV coefficients directly. Appendix C.10 presents a breakdown of the results by reduced-form and first-stage regressions.

⁵⁰ In 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.

⁵¹ This measure of effort focuses on the total number of emails sent. In app. 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 employees with higher, same, or lower pay grade.

⁵² Appendix 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, because of the precision of the coefficients, we cannot rule out considerable differences either.

⁵³ This test assumes independence of the two coefficients.

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(Δ Title) (7)
A. Posttreatment (3 months):							
[1]. Log(Manager Salary)	.150** (.074)	.130*** (.041)	.106 (.122)	-.016 (.023)	-.003 (.030)	.002 (.011)	.012 (.029)
[2]. Log(Peer Salary)	-.943** (.472)	-.431** (.210)	-.731** (.297)	.232** (.106)	.093 (.106)	.004 (.052)	.114 (.123)
B. Posttreatment (6 months):							
[3]. Log(Manager Salary)	.112 (.091)	.073 (.056)	.233*** (.087)	.004 (.030)	.012 (.030)	.029 (.048)	.010 (.029)
[4]. Log(Peer Salary)	-1.200** (.591)	-.437** (.209)	-1.064*** (.347)	.213* (.127)	-.007 (.157)	-.212 (.176)	.120 (.123)
C. Pretreatment (falsification):							
Log(Manager Salary)	.001 (.114)	-.101 (.071)	.063 (.160)	.063 (.160)	.029 (.029)	.002** (.001)	.009 (.010)
Log(Peer Salary)	-.205 (.542)	-.184 (.289)	-.191 (.412)	-.212 (.289)	.212 (.163)	-.001 (.005)	-.071* (.040)
p -value H_0 :							
[1] = [2]	.026	.007	.000	.015	.398	.963	.424
[1] = [3]	.748	.407	.400	.603	.740	.577	.980
[2] = [4]	.733	.984	.466	.913	.599	.239	.970
Cragg-Donald F statistic	29.8	204.0	98.2	203.5	203.4	203.6	203.3
Mean outcome	5.98	35.57	.48	.05	.09	.92	.10
Standard deviation outcome	1.88	44.93	.23	.21	.28	.70	.30
Observations	602	2,060	791	2,060	2,060	2,060	2,060

NOTE.—Standard errors (in parentheses) are clustered at the position level. Each column presents results for three sets of IV regressions, following the specification from sec. IV.A: in panel A, the dependent variable is the average behavior during the 90 days after the completion of the survey; in panel B, it is the average behavior during the 180 days after the completion of the survey; in panel C, it is the average behavior during the 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. $P(\text{Left})$, $P(\text{Transfer})$, and $P(\Delta\text{Title})$ 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 posttreatment (pretreatment) period. “Cragg-Donald F statistic,” “Mean outcome,” and “Standard deviation outcome” correspond to the regressions in panel A. In cols. 1–3 and 6, “Mean outcome” and “Standard deviation outcome” correspond to the values prior to taking the logarithm function. In col. 4, the pretreatment coefficient cannot be estimated because by construction there is no variation in the pretreatment outcome. Column 1 corresponds to the subsample of employees in the headquarters offices, and col. 3 corresponds to the subsample of employees with sales roles.

* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

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 thus statistically insignificant ($p = .383$).

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 $\log(\text{Peer Salary})$ that is negative (-0.943) and statistically significant ($p = .045$).⁵⁴ 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 , $p = .041$) and sales performance (-0.731 , $p = .014$).

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 $p = .026$ for hours worked (col. 1), $p = .007$ for emails sent (col. 2), and $p < .001$ for sales performance (col. 3).⁵⁵

One unique aspect of our setting is that subjects are in a continuing contract with the 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 within 90 days following 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 ($p = .029$). 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

⁵⁴ The manager coefficients are substantially more precisely estimated than the corresponding peer coefficients; e.g., in col. 1, 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, and (ii) the learning rate was higher for manager salary than for peer salary.

⁵⁵ The fact that the coefficients on $\log(\text{Peer Salary})$ are larger in absolute value than the coefficients on $\log(\text{Manager Salary})$ does not imply that horizontal comparisons are more consequential, because there is much more variation in $\log(\text{Manager Salary})$ than in $\log(\text{Peer Salary})$.

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 (col. 6) and an indicator variable showing whether the individual changed positions (col. 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.

C. *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, Wright, and Yogo 2002). Given the strong reaction to the information documented in section III.C, this information 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, Wright, and Yogo (2002); it takes the values of 29.8, 204.0, 98.2, 203.5, 203.4, 203.6, and 203.3, respectively, in columns 1–7 of table 2.

In the baseline results, we estimate the effects on behavior during the 90 days following 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 panel B of table 2. A comparison between panels A and B indicates 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 get 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- 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 II.B, we showed that, consistent with successful random assignment, the observable characteristics are balanced across the four treatment groups. Panel C of table 2 presents a more direct check. Those coefficients are estimated in a regression with the pretreatment behavior (i.e., the average during the months before the survey) instead of posttreatment 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 possibly have 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 posttreatment coefficient in column 1 of table 2 is positive (0.150, $p = .042$), while the corresponding pretreatment coefficient is close to zero (0.001) and statistically insignificant.

We conducted a number of additional robustness checks, reported in appendix C. 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 specification is also common in the literature on relative income concerns (see, e.g., Senik 2004; Clark, Frijters, and Shields 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 upward are the mirror image of updating beliefs downward. However, these effects could be asymmetric in practice; for example, 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.⁵⁶ Appendix C.15 presents the results from an econometric specification that allows for this type of asymmetric response. 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, because of power limitations, we cannot rule out small or moderate asymmetries either. Likewise, we do not find any evidence of this type of asymmetry (upward vs. downward revisions) for the horizontal comparisons.

⁵⁶ For example, employees may have more flexibility to adjust their effort upward in response to good news about the manager pay (e.g., by working extra hours) than to adjust their effort downward in response to bad news.

For the horizontal comparisons, a second type of asymmetry may arise. Indeed, evidence from related studies supports this type of asymmetry: while retention may be 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; Breza, Kaur, and Shamdasani 2018; Dube, Giuliano, and Leonard 2019). Appendix C.15 presents the results using an econometric specification 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.⁵⁷

D. 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 II.A 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 posttreatment questions included in the survey. All coefficients are estimated with the same IV specification from table 2.⁵⁸

The dependent variable in column 1 of table 3 is the (log) expected future salary 1 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 $\log(\text{Manager Salary})$ 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 5 years instead of just 1 year. According to our survey data, the average employee expects a probability

⁵⁷ The statistical power available to conduct this type of analysis is limited, so one should not conclude that the effects are perfectly symmetrical. However, the results do indicate that the baseline (symmetric) specification constitutes a reasonable approximation.

⁵⁸ The only difference is that when using survey outcomes, we do not observe pretreatment outcomes so we cannot use them for falsification tests or include them as control variables. To compensate for this lack of pretreatment controls, we include some additional control variables: dummies for sales role, pay band, unit, and position title.

TABLE 3
EFFECTS OF SALARY PERCEPTIONS ON SURVEY OUTCOMES

	LOG(E [Future Salary])		SATISFACTION		INEQUALITY TOLERANCE (5)
	+1 Year (1)	+5 Years (2)	With Pay (3)	With Job (4)	
[1]. Log(Manager Salary)	.025 (.025)	.166*** (.055)	-.015 (.125)	-.086 (.102)	-.008 (.075)
[2]. Log(Peer Salary)	.071 (.090)	.280 (.176)	-.762* (.433)	-.444 (.491)	-.373* (.216)
<i>p</i> -value H_0 :					
[1] = [2]	.595	.532	.084	.433	.135
Cragg-Donald <i>F</i> statistic	253.5	255.3	253.6	254.3	254.3
Mean dependent variable	2.58	3.22	2.79	3.60	1.80
Standard deviation					
dependent variable	.51	.59	.92	.78	.57
Observations	2,033	2,026	2,030	2,027	2,027

NOTE.—Standard errors (in parentheses) are clustered at the position level. Each column presents results for a different IV regression, following the specification described in sec. IV.A. “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 1 and 5 years in the future. “Satisfaction with Pay” and “Satisfaction with Job” are measures in a five-point scale from very dissatisfied (1) to very satisfied (5). “Inequality Tolerance” measures tolerance for pay inequality in a three-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.

* Significant at 10%.

*** Significant at 1%.

of 55.8% of being promoted to the manager’s position within the following 5 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 2 is positive (0.166), precisely estimated, and highly statistically significant ($p = .003$). A 10% increase in perceived manager salary increases the expected salary in 5 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) as well. These coefficients are also positive but less precisely estimated and thus statistically insignificant: 0.071 ($p = .431$) with respect to salary 1 year in the future (col. 1) and 0.280 ($p = .111$) with respect to salary 5 years in the future (col. 2). This evidence would also be consistent with the career concerns channel in the model of section II.A, according to which employees expect to use the information on peer salary as a bargaining chip.⁵⁹ Since

⁵⁹ This evidence is also consistent with the “tunnel effect” from Senik (2004).

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 II.A 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 they 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 5 years.⁶⁰ 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 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%.⁶¹

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 1 and 5 years, respectively. Column 1 indicates that the perceived manager salary does not affect salary expectations 1 year ahead, regardless of the distance to the manager. Column 2 indicates that the manager salary increases the employees' own salary expectations 5 years ahead but only if the employee can aspire to the managerial position: the coefficient on $\log(\text{Manager Salary})$ is positive (0.204) and statistically significant ($p = .001$) for the group Closer but smaller (0.086) and statistically insignificant ($p = .349$) for the group Farther. We find similar heterogeneity for the

⁶⁰ When compared with some objective data on career progression at this firm, these perceptions seem reasonably calibrated; for details, see app. C.3.

⁶¹ The 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 ($p < .001$).

TABLE 4
EFFECTS OF PERCEIVED MANAGER SALARY BY DISTANCE TO MANAGER

	LOG(<i>E</i> [Future Salary])		EFFORT AND PERFORMANCE		
	+1 Year (1)	+5 Years (2)	Log(Hours) (3)	Log(Emails) (4)	Log(Sales) (5)
Specification 1 (by number of promotions):					
Log(Manager Salary):					
[1]. Closer	.041 (.030)	.204*** (.059)	.212** (.099)	.170*** (.052)	.195 (.131)
[2]. Farther	-.008 (.033)	.086 (.092)	-.074 (.093)	.019 (.104)	.033 (.285)
Specification 2 (by promotion probability):					
Log(Manager Salary):					
[3]. Closer	.008 (.036)	.200*** (.059)	.431* (.226)	.185*** (.061)	.124 (.144)
[4]. Farther	.057 (.038)	.134 (.096)	-.016 (.135)	.068 (.062)	.289 (.288)
<i>p</i> -value H_0 :					
[1] = [2]	.216	.229	.040	.243	.657
[3] = [4]	.322	.560	.170	.212	.691
Observations	2,033	2,026	602	2,060	791

NOTE.—Standard errors (in parentheses) are 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 four or fewer promotions ahead. In specification 2, “Closer” indicates a probability of reaching the managerial position within 5 years of 40% or greater. The regressions in cols. 1 and 2 follow the same specification used in cols. 1 and 2 of table 3 (see note therein for more details), and the regressions from cols. 3–5 follow the same specifications used in cols. 1–3 of table 2 (see note 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*[Future Salary]” corresponds to the expected salary 1 and 5 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.

* Significant at 10%.

** Significant at 5%.

*** Significant at 1%.

effects on effort and performance. Column 3, corresponding to hours worked, shows that the coefficient on log(Manager Salary) is positive (0.212) and statistically significant ($p = .033$) for the group Closer but negative (-0.074) and statistically insignificant ($p = .424$) for the group Farther. Column 4, corresponding to emails sent, shows a coefficient on log(Manager Salary) that is positive (0.170) and statistically significant ($p = .001$) for the group Closer but smaller (0.019) and statistically insignificant ($p = .856$) for the group Farther. In column 5, corresponding to the sales outcome, the point estimate is positive (0.195) although borderline insignificant ($p = .137$) for the group Closer and close to zero (0.033) and statistically insignificant ($p = .907$) 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, 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.

E. 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. Social preferences might still be 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 $\log(\text{Manager Salary})$ on pay satisfaction is close to zero (-0.015), statistically insignificant ($p = .906$), 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 ($p = .399$ and $.920$), 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 ($p = .078$). 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 $p = .084$. The effect of peer salary on

job satisfaction is statistically insignificant ($p = .366$) but still negative (-0.444 , from col. 4) and similar in magnitude and statistically indistinguishable from the corresponding effect on pay satisfaction (-0.762 , from col. 3). The peer salary coefficient from column 5 is negative (-0.373) and statistically significant ($p = .084$). 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.⁶² 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 maintains that individuals tend to compare themselves not with everyone else but with other people in specific reference groups (Clark and Senik 2010). In our context, employees may find it natural to compare their own salaries with the salaries of other employees in the same position but may not feel inclined to compare themselves with 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 perceived as unfair.

A third interpretation is also based on fairness concerns. Employees may be intolerant of the horizontal salary differences because they perceive those differences as arising largely out of nonmeritocratic factors, such as luck and office politics. Indeed, the findings from Breza, Kaur, and Shamdasani (2018) suggest that small horizontal differences in pay can be justified if they arise because of large differences in productivity. By contrast, employees may find it much 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, Dubra, and Lagomarsino 2016).

While we do not have direct evidence on whether employees perceive horizontal salary differences to be meritocratic, we do have some

⁶² While 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, the discussion of other potential mechanisms is reported in app. C.16.

suggestive evidence that nonmeritocratic factors may play a significant role in this context. Several factors can be at play when determining horizontal salary differences. While rules govern 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.⁶³ In principle, some of these factors may be perceived as meritocratic, while other factors may be perceived as nonmeritocratic. For example, some employees may perceive that some of their peers are paid more because 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 finding as evidence that employees may perceive horizontal differences to be largely nonmeritocratic.

V. Discussion

A. 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 in which employees are regularly promoted to higher positions. Second, employees should be employed in long-term contracts with a well-defined base salary. Third, pay-setting practices, such as transparency, performance pay, and inequality, should be representative of most large firms. All of 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 also in any of those other firms, from a tech company in Silicon Valley and a law firm in Jakarta to a hospital in 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 be large enough to be able to recruit thousands of employees for the experiment. The test demanded that we could use administrative records to track the behavior of the employees, such as their effort and performance, and

⁶³ For a detailed description of how salaries are determined at this organization, see app. C.17.

observe sensitive information, such as the employees' salaries, the organization structure, and the network between coworkers. 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. For these reasons, our firm was the ideal setting for the test.

As discussed in section II.B, 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, which does not mean that we expect the results from this study to be identical if conducted in a different firm. On the contrary, according to the economic theories for which we test, 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 among the employees of this firm, the motivating effects were significant only for those who could aspire to attain the managerial position. In other firms, opportunities to climb the ladder may be negligible, and so should the motivating effects of vertical comparisons. For example, the drivers for ride-sharing applications 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 this 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 thus 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).⁶⁴ Since we are the first researchers to study vertical comparisons, more work 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, Kaur, and Shamdasani (2018) conducted their experiment with manufacturing workers from a small Indian firm, Cohn et al. (2014) conducted their experiment with sales representatives from a small German firm, Dube, Giuliano, and Leonard (2019) leveraged 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 settings involving 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 overestimate 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.⁶⁵

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.⁶⁶ As an experimental context, attrition is not an issue because the outcomes can be tracked with administrative records. In terms of the 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 context has no concerns regarding scalability. Given the nature of our

⁶⁴ The similarity of social preferences across countries is also consistent with the view held by anthropologists that social preferences are human universals (Brown 1991).

⁶⁵ Appendix C.5 presents the heterogeneity in misperceptions and willingness to pay, app. C.8 presents the heterogeneity in learning, and app. C.13 presents some heterogeneity in the effects of perceptions on behavior.

⁶⁶ For details, see app. C.3.

intervention—providing simple information to the subjects—scaling the intervention to the universe of employees should be straightforward, and also, owing to the nature of the intervention, 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 that 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 that our research design will be applied in other firms and other countries to shed light on the mediating factors and to provide further support for the underlying economic models.

B. 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 III.A, 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 this additional insight. 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—that is, the average salaries by position and unit.⁶⁷ 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. The finding that employees report a desire for more pay transparency has been documented in other firms also; for example, 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 significantly depends on how the information is disclosed. In a second scenario, we offer employees to replace the status quo policy on 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

⁶⁷ 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.” For more details, see app. C.18.

employees can look up your salary as well. Employees showed little support 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 highly value the salary information, 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 IV 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 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 underestimated. 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 IV 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 also relate to some of the debates around pay transparency. The evidence of significant information frictions presented in section III.A indicates that policies that mandate pay disclosure may go a long way toward 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 improving 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.⁶⁸ 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 V.A, 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, some specific features of our transparency intervention may be important to keep in mind as well. For 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 caveat is important to keep in mind because some pay transparency policies are aimed precisely at disclosing the salaries of the top executives.

VI. 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

⁶⁸ It 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 as well. Also, our results indicate that pay transparency may still be useful for inducing horizontal pay compression, such as paying employees the same within a given position or paying men and women equally for the same job.

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, Jensen, and Murphy 1988).

Our findings also have implications for pay transparency. Our results suggest that most employees would be better off if their employers were more transparent about pay; 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.⁶⁹ Our findings suggest that these policies may be 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 warrant consideration 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, George A. 1982. "Labor Contracts as Partial Gift Exchange." *Q.J.E.* 97 (4): 543–69.
- Akerlof, George A., and Janet L. Yellen. 1990. "The Fair Wage-Effort Hypothesis and Unemployment." *Q.J.E.* 105 (2): 255–83.
- Allcott, Hunt, and Judd B. Kessler. 2019. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons." *American Econ. J.: Appl. Econ.* 11 (1): 236–76.
- Al-Ubaydli, Omar, and John A. List. 2013. "On the Generalizability of Experimental Results in Economics." In *Handbook of Experimental Economic Methodology*, edited by Guillaume R. Fréchet and Andrew Schotter, 420–62. Oxford: Oxford Univ. Press.
- Al-Ubaydli, Omar, John A. List, Danielle LoRe, and Dana Suskind. 2017. "Scaling for Economists: Lessons from the Non-adherence Problem in the Medical Literature." *J. Econ. Perspectives* 31 (4): 125–44.
- Andersen, Steffen, Glenn W. Harrison, Morten Igel Lau, and E. Elisabet Rutström. 2006. "Elicitation Using Multiple Price List Formats." *Experimental Econ.* 9 (4): 383–405.
- Angulo, Ana M., José M. Gil, and L. Tamburo. 2005. "Food Safety and Consumers' Willingness to Pay for Labelled Beef in Spain." *J. Food Products Marketing* 11 (3): 89–105.

⁶⁹ For example, policy makers promote transparency policies, such as disclosure of CEO pay (Faleye et al. 2013; Mas 2016; Mueller, Ouimet, and Simintzi 2017), with the hope of putting pressure on firms to reduce pay inequality.

- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data." *Q.J.E.* 109 (4): 881–919.
- Baker, George, Michael Jensen, and Kevin Murphy. 1988. "Compensation and Incentives: Practice vs. Theory." *J. Finance* 43 (3): 593–616.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak. 1964. "Measuring Utility by a Single-Response Sequential Method." *Behavioral Sci.* 9 (3): 226–32.
- Bottan, Nicolas L., and Ricardo Perez-Truglia. 2020. "Choosing Your Pond: Location Choices and Relative Income." *Rev. Econ. and Statis.*, forthcoming.
- Bracha, Anat, Uri Gneezy, and George Loewenstein. 2015. "Relative Pay and Labor Supply." *J. Labor Econ.* 33 (2): 297–315.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2018. "The Morale Effects of Pay Inequality." *Q.J.E.* 133 (2): 611–63.
- Brown, Donald E. 1991. *Human Universals*. Philadelphia: Temple Univ. Press.
- Burchell, Brendan, and Dana Yagil. 1997. "Socioeconomic and Political Initiators of Pay Comparison." *Work, Employment, and Soc.* 11 (4): 737–48.
- Caldwell, Sydnee, and Nikolaj Harmon. 2018. "Outside Options, Bargaining and Wages: Evidence from Coworker Networks." Working paper, Massachusetts Inst. Tech., Cambridge, MA.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *A.E.R.* 102 (6): 2981–3003.
- Carson, Richard T., and Theodore Groves. 2007. "Incentive and Informational Properties of Preference Questions." *Environmental and Resource Econ.* 37 (1): 181–210.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Econ. J.: Macroeconomics* 9 (3): 1–35.
- Charness, Gary, Uri Gneezy, and Brianna Halladay. 2016. "Experimental Methods: Pay One or Pay All." *J. Econ. Behavior and Org.* 131:141–50.
- Clark, Andrew E., Paul Frijters, and Michael A. Shields. 2008. "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *J. Econ. Literature* 46 (1): 95–144.
- Clark, Andrew E., and Andrew J. Oswald. 1996. "Satisfaction and Comparison Income." *J. Public Econ.* 61 (3): 359–81.
- Clark, Andrew E., and Claudia Senik. 2010. "Who Compares to Whom? The Anatomy of Income Comparisons in Europe." *Econ. J.* 120 (544): 573–94.
- CNBC. 2017. "A Brief History of the 8-Hour Workday, Which Changed How Americans Work." CNBC, May 3.
- Cohn, Alain, Ernst Fehr, Benedict Herrmann, and Frédéric Schneider. 2014. "Social Comparison and Effort Provision: Evidence from a Field Experiment." *J. European Econ. Assoc.* 12:877–98.
- Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment." *J. Public Econ.* 98:100–112.
- Cullen, Zoë, and Bobak Pakzad-Hurson. 2016. "Equilibrium Effects of Pay Transparency in a Simple Labor Market." Working Paper no. 28903, NBER, Cambridge, MA.
- Cullen, Zoë, and Ricardo Perez-Truglia. 2018. "The Salary Taboo: Privacy Norms and the Diffusion of Information." Working Paper no. 25145, NBER, Cambridge, MA.
- Danz, David, Lise Vesterlund, and Alistair J. Wilson. 2020. "Belief Elicitation: Limiting Truth Telling with Information on Incentives." Working Paper no. 27327, NBER, Cambridge, MA.

- Day, Nancy E. 2007. "An Investigation into Pay Communication: Is Ignorance Bliss?" *Personnel Rev.* 36 (5): 739–62.
- Dellavigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2019. "Estimating Social Preferences and Gift Exchange with a Piece-Rate Design." Working Paper no. 22043, NBER, Cambridge, MA.
- Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The Economics of Career Concerns, Part I: Comparing Information Structures." *Rev. Econ. Studies* 66 (1): 183–98.
- Di Tella, Rafael, Juan Dubra, and Alejandro L. Lagomarsino. 2016. "Meet the Oligarchs: Business Legitimacy, State Capacity and Taxation." Working Paper no. 22934, NBER, Cambridge, MA.
- Dube, Arinjadrit, Laura Giuliano, and Jonathan Leonard. 2019. "Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior." *A.E.R.* 109 (2): 620–63.
- Dvorak, Phred. 2007. "Limits on Executive Pay: Easy to Set, Hard to Keep." *Wall Street J.*, April 9.
- Edwards, Matthew A. 2005. "The Law and Social Norms of Pay Secrecy." *Berkeley J. Employment and Labor Law* 26:41–63.
- Faleye, Olubunmi, Ebru Reis, and Anand Venkateswaran. 2013. "The Determinants and Effects of CEO-Employee Pay Ratios." *J. Banking and Finance* 37 (8): 3258–72.
- Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *Q.J.E.* 133 (4): 1645–92.
- Frank, Robert H. 1984. "Are Workers Paid Their Marginal Products?" *A.E.R.* 74 (4): 549–71.
- Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2018. "Expectations with Endogenous Information Acquisition: An Experimental Investigation." *Rev. Econ. and Statis.*, forthcoming.
- Gee, Kelsey. 2017. "Ask Me How Much Money I Make: Pay Gets More Transparent." *Wall Street J.*, October 26.
- Gibbons, Robert, and Kevin J. Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *J.P.E.* 100 (3): 468–505.
- Gibbons, Robert, and Michael Waldman. 1999a. "Careers in Organizations: Theory and Evidence." In *Handbook of Labor Economics*, vol. 3B, edited by Orley C. Ashenfelter and David Card, 2373–437. Amsterdam: North-Holland.
- . 1999b. "A Theory of Wage and Promotion Dynamics inside Firms." *Q.J.E.* 114 (4): 1321–58.
- Glassdoor. 2016. "Global Salary Transparency Survey." Report, Glassdoor, San Francisco.
- Gneezy, Uri, and John A. List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica* 74 (5): 1365–84.
- Godechot, Olivier, and Claudia Senik. 2015. "Wage Comparisons in and Out of the Firm: Evidence from a Matched Employer-Employee French Database." *J. Econ. Behavior and Org.* 117:395–410.
- Harris, Milton, and Bengt Holmstrom. 1982. "A Theory of Wage Dynamics." *Rev. Econ. Studies* 49 (3): 315–33.
- Harrison, Glenn W., and John A. List. 2004. "Field Experiments." *J. Econ. Literature* 42 (4): 1009–55.
- Hegewisch, Ariane, Claudia Williams, and Robert Drago. 2011. "Pay Secrecy and Wage Discrimination." Report, Inst. Women's Policy Res., Washington, DC.

- Hoff, Peter D. 2009. *A First Course in Bayesian Statistical Methods*. Dordrecht: Springer.
- Holmstrom, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Rev. Econ. Studies* 66 (1): 169–82.
- Huet-Vaughn, Emiliano. 2017. "Do Social Comparisons Motivate Workers?" Working paper, Middlebury Coll., Middlebury, VT.
- Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2): 467–75.
- Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *J.P.E.* 87 (5): 972–90.
- Khattak, Asad J., Youngbin Yim, and Linda Stalker Prokopy. 2003. "Willingness to Pay for Travel Information." *Transportation Res. C: Emerging Tech.* 11 (2): 137–59.
- Kiatpongsan, Sorapop, and Michael I. Norton. 2014. "How Much (More) Should CEOs Make? A Universal Desire for More Equal Pay." *Perspectives Psychological Sci.* 9 (6): 587–93.
- Lawler, Edward E. 1965. "Managers' Perceptions of Their Subordinates' Pay and of Their Superiors' Pay." *Personnel Psychology* 18 (4): 413–22.
- . 2003. "Pay Practices in Fortune 1000 Corporations." *WorldatWork J.* 12 (4): 45–54.
- Lazear, Edward P. 1989. "Pay Equality and Industrial Politics." *J.P.E.* 97 (3): 561–80.
- Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *J.P.E.* 89 (5): 841–64.
- Levitt, Steven D., and John A. List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *J. Econ. Perspectives* 21 (2): 153–74.
- List, John A. 2020. "Non Est Disputandum de Generalizability? A Glimpse into the External Validity Trial." Working Paper no. 27535, NBER, Cambridge, MA.
- Mas, Alexandre. 2016. "Does Disclosure Affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act." Working paper, Princeton Univ., Princeton, NJ.
- . 2017. "Does Transparency Lead to Pay Compression?" *J.P.E.* 125 (5): 1683–721.
- McKenzie, David. 2012. "Beyond Baseline and Follow-Up: The Case for More T in Experiments." *J. Development Econ.* 99 (2): 210–21.
- Mueller, Holger M., Paige P. Ouimet, and Elena Simintzi. 2017. "Within-Firm Pay Inequality." *Rev. Financial Studies* 30 (10): 3605–35.
- PayScale. 2018. "2018 Compensation Best Practices Report." Report, PayScale, Seattle.
- Perez-Truglia, Ricardo. 2020. "The Effects of Income Transparency on Well-being: Evidence from a Natural Experiment." *A.E.R.* 110:1019–54.
- Romer, David. 1984. "The Theory of Social Custom: A Modification and Some Extensions." *Q.J.E.* 99 (4): 717–27.
- Rosen, Sherwin. 1986. "Prizes and Incentives in Elimination Tournaments." *A.E.R.* 76 (4): 701–15.
- Senik, Claudia. 2004. "When Information Dominates Comparison: Learning from Russian Subjective Panel Data." *J. Public Econ.* 88 (9): 2099–123.
- Shogren, Jason F., Sungwon Cho, Cannon Koo, John List, Changwon Park, Pablo Polo, and Robert Wilhelm. 2001. "Auction Mechanisms and the Measurement of WTP and WTA." *Resource and Energy Econ.* 23:97–109.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2019. "Firming Up Inequality." *Q.J.E.* 134 (1): 1–50.

- Stock, James H., Jonathan H. Wright, and Motohiro Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *J. Bus. and Econ. Statis.* 20 (4): 518–29.
- Trachtman, Richard. 1999. "The Money Taboo: Its Effects in Everyday Life and in the Practice of Psychotherapy." *Clinical Soc. Work J.* 27 (3): 275–88.
- Wen, Qi, Peter A. Gloor, Andrea F. Colladon, Praful Tickoo, and Tushar Joshi. 2020. "Finding Top Performers through Email Patterns Analysis." *J. Information Sci.* 46 (4): 508–27.