The Salary Taboo

Privacy Norms and the Diffusion of Information

Zoë Cullen — *Harvard Business School*

Ricardo Perez-Truglia* — *University of California, Los Angeles*

First Draft: October 2018.

Abstract

The diffusion of salary information has important implications for labor markets, such as for wage discrimination policies and collective bargaining. Despite the widespread view that transmission of salary information is imperfect and unequal, there is little direct evidence on the magnitude and sources of these frictions. We conduct a field experiment with 752 employees at a multibillion-dollar corporation to address these questions. We provide evidence of significant frictions in how employees search for and share salary information and suggestive evidence that these frictions are due to privacy norms. We do not find any significant differences in information frictions between female and male employees.

JEL Classification: D83, D84, D91, C93, J16, J31, M12.

Keywords: information diffusion, salary, privacy, inequality, transparency, gender.

*Cullen: zcullen@hbs.edu, Rock Center 312, Boston, MA 02163. Perez-Truglia: ricardo.truglia@anderson.ucla.edu, 110 Westwood Plaza, Los Angeles, CA 90095. We are thankful for comments from several colleagues and seminar discussants. The collaborating firm provided financial support for the research being conducted. Additionally, Zoë Cullen was a full-time, salaried employee at that firm while the research was being conducted. We thank Katherine Fang for excellent research assistance. This project was reviewed and approved by the Institutional Review Board at University of California Los Angeles (IRB#18-001467).

1 Introduction

Most employers provide limited information about salaries. Thus, employees' knowledge about salaries depends largely on their ability to communicate with each other. However, there is a widespread belief that the diffusion of salary information is imperfect and unequal. For example, most employees do not discuss salaries with their coworkers, despite wanting to be better informed about peers' salaries (Glassdoor, 2016; PayScale, 2018). These information frictions are sometimes attributed to firm efforts that discourage employees from discussing salaries (Gely and Bierman, 2003; Hegewisch et al., 2011).[1] Others argue that the frictions stem from a "salary taboo": a social norm around salary privacy that discourages coworkers from revealing or inquiring about salary information (Trachtman, 1999; Edwards, 2005).

These information frictions are important, because they have implications for a broad range of labor market phenomena. For example, information frictions can facilitate workplace discrimination, increase employers' market power (Danziger and Katz, 1997; Cullen and Pakzad-Hurson, 2017), and hinder collective bargaining and unionization (Corbett, 2002). These supposed information frictions also have inspired several policies, such as those that punish employers when they retaliate against employees who discuss wages with each other (Pender, 2017; Siniscalco et al., 2017). Despite these important implications, there is little *direct* evidence on the diffusion of salary information. We use a field experiment to provide novel evidence on how individuals search for and share salary information and on the role that privacy concerns play in these decisions.

Employees can benefit from information about coworkers' salaries in several scenarios, such as negotiating salary, switching managers, or searching for new jobs. This information has a cost, though: employees must spend time and energy to search for it, and they may face costs for inquiring about sensitive data.[2] We design a novel field experiment to study these costs and benefits.

To measure employees' *willingness to search for information,* we generate an exogenous shock to the benefits of being informed about salaries by allowing employees to partake in a game. In the game, employees guess the average salary of a random sample of five of their peers (e.g., a bank teller guesses the average salary of five other tellers from the same branch). Employees whose guesses fall within 5% of the true average salary receive

[1] There is no consensus for why firms prefer pay secrecy or whether it is in the firm's best interest. Some argue that firms use secrecy to undermine collective bargaining (Bierman and Gely, 2004), reduce manipulative behavior (Brickley et al., 2007), or avoid the diffusion of information about outside offers (Danziger and Katz, 1997).

[2] Websites that aggregate salary information, such as Glassdoor and PayScale, can lower the social costs of searching. However, these websites offer an imperfect substitute from searching for information within the organization.

a monetary reward. After providing their initial guesses, which they must do immediately, respondents are offered the opportunity to acquire an extra week to search for information in the wild and improve their guesses. We elicit the probability of winning the game with and without the additional week, using self-reported and incentive-compatible methods. The degree to which employees expect the extra week to increase their probability of winning the game measures their willingness to search for information in the wild.

We measure the *gross benefits from information* by eliciting the willingness to pay for an imperfect but informative signal about the average peer salary, using an incentive-compatible method. In other words, we provide subjects with the opportunity to acquire readily available information from the experimenter instead of searching for information in the wild. Last, we measure the *willingness to share information with others*. We offer respondents an opportunity to reveal their own salaries to five peers, and we use an incentive-compatible method to elicit the willingness to pay to reveal this information (for subjects who prefer to share the information) or the willingness to pay to conceal this information (for subjects who prefer to conceal the information).

We cross-randomize two key features of the survey. The first treatment arm allows us to test differences in the diffusion of salary information, relative to other important career information. For this, we randomize subjects into two versions of the survey: salary and seniority. The survey types are identical, except that one asks about the average *seniority* of peers instead of the average *salary*. Just like information about peer salary, information about peer seniority can be useful to make important career choices, such as whether to ask for a promotion or search for another job. However, employees may face higher frictions when searching for and sharing information about salary, compared to seniority, for example, because of the salary taboo.

The second treatment arm aims to test the rational inattention hypothesis, according to which individuals search for and acquire new information when they stand to gain from it (Woodford, 2001; Sims, 2003; Mankiw and Reis, 2002; Reis, 2006). We randomize the size of the rewards of the guessing game using five different values from \$13 to \$63 (these and all other monetary values reported in this paper are expressed in United States dollars using PPP-adjusted exchange rates from February 2018). This randomization generates exogenous variation in the benefits of being informed. We test two predictions of the rational inattention model: higher rewards should increase the willingness to search for information in the wild and the willingness to acquire readily available information.

We conduct a field experiment with a sample of 752 employees from a large commercial bank (hereafter referred to as *the firm*) with thousands of employees, millions of customers, and billions of dollars in revenues. The firm is typical in some relevant respects. The firm

does not have open salary policies and discourages employees from discussing salaries with each other. Most employees report that they have limited information about salaries and would prefer the firm to be more transparent. There seems to be a social norm against asking coworkers about their salaries, and employees rarely discuss salaries with their coworkers. A number of studies show that these features are common in firms from several countries, including the United States (Trachtman, 1999; Edwards, 2005; Hegewisch et al., 2011; Glassdoor, 2016; PayScale, 2018).

We find that employees have imperfect information about the salaries of their peers: the mean absolute error of the guesses reported in the game is 16%. Indeed, this level of misperception is what we would expect if employees have access only to information about their own salaries. Although employees are overconfident in their guesses, they are aware that their accuracy (i.e., the probability of guessing within 5% of the truth) is far from perfect.

We provide evidence that misperceptions are partly due to search costs. When presented with financial incentives to do so, most individuals are willing to search for information in the wild. When given an extra week to gather information, the average respondent expects to increase the probability of winning the guessing game by 23 percentage points. The evidence reveals that search costs are unequal: when provided with the additional week, some employees expect to search for information in the wild, but other employees do not expect to search. And consistent with the rational inattention hypotheses, employees who are randomly assigned to higher game rewards expect to search more intensively than employees who are assigned to lower game rewards.[3]

We find that, much like the search costs, the gross benefits of the salary information are significant and unevenly distributed. The median willingness to pay for the readily available signal of peer salary is about $13. Consistent with rational inattention, this value is higher for individuals who are assigned to a higher game reward and thus stand to gain more from the information. Employees in the bottom half of the distribution, who are willing to pay less than $13 for the information, seem to be misinformed mainly due to a lack of interest. On the other hand, the remaining half of subjects highly value the information: their willingness to pay for the signal has a median of $130 and a mean of $369. These high valuations suggest that these employees do not search for information in the wild because of information frictions.

We find that individuals also face significant frictions when sharing information with others. The willingness to reveal one's salary to coworkers is both significant and heterogeneous. Whereas a minority of employees (20%) prefer to share personal salary information with their

[3]These findings are consistent with growing evidence of rational inattention in contexts such as hiring (Bartoš, Bauer, Chytilová, and Matějka, 2016), inflation expectations (Cavallo, Cruces, and Perez-Truglia, 2017), and housing expectations (Fuster, Perez-Truglia, and Zafar, 2018).

peers, most (80%) prefer to conceal this information. Moreover, this preference for privacy can be strong: some employees would reveal their salaries for a small sum of money, but roughly half would not be persuaded to reveal the information to five peers even for $125.

The preference for privacy is consistent with a salary taboo. Individuals are afraid to ask coworkers about their salaries, because they understand that most coworkers prefer to keep their salary information private. Indeed, this interpretation is consistent with our subjective data. Most respondents report that it is socially unacceptable to ask coworkers about their salaries and that they feel uncomfortable doing so. Moreover, 89% of respondents believe that if they ask coworkers about their salaries, they will get asked about their own salaries. Thus, employees may be afraid to ask coworkers about their salaries because that may force them to reveal their own salaries, which they dislike.

We find that employees are better informed about peer seniority than about peer salaries. When guessing salaries, employees are as accurate as they would be if they just reported their own salaries. This finding indicates that employees do not have access to information beyond their own salaries. In contrast, when guessing seniority, employees are substantially more accurate than they would be if they just reported their own seniority. This finding suggests that employees have access to other information about seniority besides their own seniority. Moreover, our evidence suggests that, at the margin, employees stand to gain more from salary information than from seniority information. This evidence suggests that the difference between salary misperceptions and seniority misperceptions are due to differences in search costs.

Although it is not the only possible interpretation, the demand for privacy is our favorite interpretation for the differences between salary misperceptions and seniority misperceptions. Two main pieces of evidence support this view. First, the revealed-preference evidence indicates that the topic of salary is substantially more sensitive than that of seniority. The average employee is willing to reveal personal salary information to a sample of five peers for $67 and willing to reveal seniority to peers for just $28. Second, the subjective data also suggest that salary is a more sensitive topic. Whereas 69% of employees find it unacceptable to ask a coworker about salary, only 6% find it unacceptable to ask about seniority; and whereas 53% of employees find it uncomfortable to ask about a coworker's salary, only 5% find it uncomfortable to ask about seniority.

We find substantial frictions in information diffusion, even though participation in the game may facilitate this diffusion. For example, the guessing game may provide an excuse to ask peers about their salaries that mitigates the fears of breaking a social norm or breaking the company's disclosure rules. Thus, if anything, our findings may underestimate the magnitude of information frictions under normal circumstances.

Our last result relates to gender differences in information frictions. This analysis is motivated by the widespread view that pay secrecy disproportionately affects women (Babcock and Laschever, 2009) and thus may be one of the factors behind the gender pay gap.[4] Consistent with this view, survey data indicate that women are less confident than men about their salary knowledge (Glassdoor, 2016; Cullen and Pakzad-Hurson, 2017). Consistent with these prior survey findings, our own data indicate that female employees are less confident than male employees about their ability to guess the salaries of their peers. However, we find that those differences in confidence do not correspond with any real differences in accuracy. If anything, female employees are slightly more accurate than their male counterparts. Moreover, we find that other gender differences are small, statistically insignificant, and precisely estimated: female and male employees are equally willing to search for information, equally willing to buy information, and equally willing to share information with peers.

Our study relates to various strands of literature. A large theoretical literature from economics and management suggests that frictions in the diffusion of salary information can have important implications for labor markets (Akerlof and Yellen, 1990; Kuhn and Gu, 1998, 1999; Ellingsen and Rosén, 2003; Michelacci and Suarez, 2006; Cullen and Pakzad-Hurson, 2017; Moellers, Normann, and Snyder, 2017). Yet, there is little *direct* evidence on the magnitude and sources of information frictions.[5] This study builds on our previous work documenting significant misperceptions of peer and manager salaries (Cullen and Perez-Truglia, 2018).[6] This study aims to understand the sources of these misperceptions, with special emphasis on the role of the salary taboo.

Our study relates to a literature on the diffusion of information in social networks. Several models explain how individuals form beliefs based on peer-to-peer communication (Bass, 1969; Ellison and Fudenberg, 1995). More recent studies measure social learning in the field (Mobius and Rosenblat, 2014). Some of these studies artificially create incentives for information diffusion. For instance, Mobius et al. (2015) recruited college students to play a "treasure hunt" game in which they earned prizes by collecting information from peers. Other studies exploit natural incentives for information diffusion. For example, Beaman et al. (2018) seeded useful information about composting and measured its diffusion in an agricultural network. These papers show evidence that, even in settings where information

[4]For example, the landmark case against gender wage discrimination in the United States, the Lilly Ledbetter Fair Pay Act of 2009, touches extensively on the topic of information diffusion.

[5]A small but growing literature shows that changes in pay transparency can affect employee behavior and employee satisfaction, which constitutes suggestive evidence of information frictions (Card, Mas, Moretti, and Saez, 2012; Perez-Truglia, 2015; Mas, 2016, 2017; Breza, Kaur, and Shamdasani, 2018; Cullen and Pakzad-Hurson, 2017; Cullen and Perez-Truglia, 2018).

[6]This earlier work provides evidence that employees have misperceptions about the salaries of their peers and managers, that they are sometimes willing to pay significant amounts to acquire accurate information, and that they do not share information with their peers when it is provided.

is mutually beneficial, its diffusion is highly imperfect. Our contribution to this literature is twofold. First, we contribute a new method to measure the willingness to search for information and the willingness to share information with others. Second, we explore the role of privacy norms for the diffusion of information.

Our paper adds to the literature on the economics of privacy (Acquisti et al., 2016). For example, Goldfarb and Tucker (2012) show that, even in anonymous internet surveys, some respondents refuse to reveal information about their incomes and demographics. Athey et al. (2017) and Adjerid et al. (2013) study the demand for privacy in the crypto-currency market. They show that even individuals who report that they highly value privacy are willing to give away sensitive information for small incentives. We contribute to this literature by measuring preferences for privacy in a context with high stakes (i.e., an employee's willingness to reveal personal salary information to coworkers). In contrast to those other contexts, we find a high willingness to pay for privacy. Perhaps more surprisingly, we find a large heterogeneity in preferences for privacy, with some individuals willing to pay to reveal their salary to peers rather than conceal it.

Last, this study relates to literature on wage discrimination. There is a widespread view that pay secrecy hurts minorities, because it helps employers to discriminate against them (Phillips, 2009; Colella et al., 2007). This view has led to various efforts to reduce the gender wage gap, through transparency policies (Colella et al., 2007). However, this argument often assumes that pay secrecy hinders information access more for women and minorities than for others. Our evidence does not support this assumption: women and men face similar frictions and have similar degrees of misperceptions. However, we do find that female employees are less confident than male employees about the accuracy of their beliefs.

The rest of the paper proceeds as follows. Section 2 presents the conceptual framework. Section 3 presents the survey design. Section 4 discusses the implementation details. Section 5 presents the results. The last section concludes.

2 Conceptual Framework

In this section, we provide the conceptual framework that motivates the design of the survey and aids in the interpretation of its results.

2.1 Basic Framework

Figure 1.a provides a simple framework for thinking about how employees search for information about the salaries of their peers. We assume that employees care about their coworkers'

average salary, and that their utility depends on the accuracy of this belief – for the sake of simplicity, assume this accuracy is equal to the perceived probability that individual's belief falls within 5% of the truth. The employee's accuracy is represented in the x-axis. Searching for information would allow individuals to move towards the right in the x-axis. For instance, finding new information about a peer's salary should increase the perceived accuracy by some positive number. This figure also shows the marginal benefit (MB) and marginal cost (MC) curves. The q_0 point on the x-axis represents the point at which a rational individual stops acquiring new information, exactly where the MC and MB curves intersect.

Figure 1.a perhaps depicts the most natural case, in which the MC is increasing and the MB is decreasing. To motivate the upward slope of MC, consider individuals who do not ask peers about their salaries.[7] These individuals may be more comfortable asking some peers than others. They start with the peer they feel most comfortable asking, then move to the second most comfortable, and so on. Thus, the marginal cost increases with each additional piece of information. The downward slope of the MB curve is based on the principle of diminishing marginal returns: a 1% increase in accuracy is more useful to an individual who is completely uncertain about the topic than it is to an individual who is almost certain. The assumptions that MB is decreasing and MC is increasing help simplify the explanation, but they are not crucial for the following results.[8]

2.2 Willingness to Search for Information in the Wild

Figure 1.b presents our measure of willingness to search for information. The point q_0 represents the point in an individual's search where the marginal benefit equals the marginal cost. At this point, the individual is given the unexpected opportunity to participate in a guessing game, wherein correctly guessing the average salary of peers within 5% of the truth earns a reward. The introduction of the game shifts the MB curve upwards, to MB′. For example, if the individual is risk-neutral, the MB shifts upward by an amount equal to the game reward.

Individuals who are not given any time to search for information before providing a guess stay at q_0. In this case, playing the guessing game increases their utility by an amount equal to the area of the shaded parallelogram from Figure 1.b. This area represents the willingness to forfeit the guessing game. If individuals have extra time to search for information before providing the guess, they will want to search for additional information up to point q_1, where

[7] Individuals do not need to ask their peers directly to search for information. For example, they could consult information provided by the employer, such as pay bands, or they could ask one employee about another employee's salary.

[8] A more important assumption is whether the curves intersect exactly once, which is guaranteed when the MB is decreasing and the MC is increasing. If the curves intersect multiple times, then the analysis gets a little more complicated, but the main intuitions remain.

the MB′ curve intersects the MC curve. This expected gain in accuracy, $q_1 - q_0$, measures the individual's willingness to search for information (i.e., the higher this gap, the more the individual searches in response to the game incentives).

This willingness to search can be estimated by the difference between the expected probability of winning the guessing game with and without extra time. One challenge is that individuals may not report these probabilities truthfully. Thus, as a robustness check, we provide an alternative incentive-compatible measure of the willingness to search: the shaded triangle between q_0 and q_1 from Figure 1.b, which corresponds to the willingness to pay to acquire extra time to search for information to improve one's guess. We can elicit this willingness to pay using standard incentive-compatible methods. We normalize this willingness to pay by dividing it by the reward amount. Indeed, if the individual were risk-averse, this normalized measure should equal $q_1 - q_0$.

To illustrate the first rational inattention hypothesis, Figure 1.c shows two hypothetical scenarios in which the individual is offered different reward amounts. With a lower reward, the MB curve shifts upward to MB′. With a higher reward, the MB curve shifts even further, to MB′′. When facing the higher reward, the rational individual responds by searching for more information, up to the point $q_2 > q_1$. In our research design, we provide a test of this hypothesis by randomizing the reward amount.

2.3 Willingness to Acquire Readily Available Information

Figure 1.d illustrates two potential scenarios, corresponding to the MC and MB curves with subscripts 1 and 2, respectively. These two scenarios generate the same willingness to search for information. However, they correspond to substantially different gross benefits from search. In the first scenario, corresponding to curves MB_1 and MC_1, the gross benefits from information are large, implying that employees may gain a lot from the removal of information frictions. In the alternative scenario, corresponding to curves MB_2 and MC_2, gross benefits from information are small, implying that employees have little to gain from the removal of information frictions.

Figure 1.e shows how to measure the gross benefits from information. Consider individuals who are given the opportunity to play the guessing game but no additional time to acquire information. These individuals can buy a signal related to the average peer salary and then use that signal to revise their guesses in the game. The individuals expect that, after processing the signal, the accuracy of their guesses will increase from q_0 to q_s. Unlike in the previous graphs, this q_s is exogenously given by the quality of the signal, and thus unrelated

to the point where the MC and MB curves intersect.[9] The willingness to pay for this signal equals the area of the shaded trapezoid from Figure 1.e.[10] The area of this trapezoid allows us to distinguish between the two alternative scenarios in Figure 1.d.

Moreover, the reward randomization allows us to test a second rational inattention hypothesis. Figure 1.f illustrates two hypothetical scenarios in which the individual is offered different reward amounts. With a lower reward, the MB curve shifts upward to MB′. With a higher reward, the MB curve shifts even further, to MB′′. When facing the higher reward, the rational individual should be willing to pay more for the signal. In Figure 1.f, this extra demand for the signal is equal to the area of the shaded parallelogram.

3 Survey Design

3.1 Survey Types: Salary Versus Seniority

Appendix C includes a full sample of the survey instrument. Participants are assigned with equal probability to one of two types of surveys:

- *Salary Survey*: this survey type asks about the average salary among peers. We use the standard definition of peers: other employees who share the same position title and work in the same organizational unit (Card et al., 2012; Cullen and Perez-Truglia, 2018). For instance, the peers of a teller in a bank branch could be the other tellers in the same branch. We use one specific type of salary, the monthly gross base salary, which we describe in detail in the survey. This salary excludes any additions or deductions, such as taxes, allowances, commissions, or bonuses. According to interviews with the HR department and employees who did not participate in the experiment, this salary type is the most salient for employee compensation and is typically the most relevant figure in the employee's contract. Base salary also is the total compensation amount for nearly all subjects in our sample.

- *Seniority Survey:* this survey type asks about the average seniority of peers, which is defined as the number of years elapsed since the employee joined the company.

The two types of survey instruments are identical, except that the word "salary" replaces all instances of "seniority" and the corresponding "$" units replace all instances of the "years"

[9] Under Bayesian learning, this expected increase in accuracy depends primarily on the precision of the prior belief and the precision of the signal.

[10] We can decompose the trapezoid into two parts: the bottom part (a smaller trapezoid) is the willingness to pay for knowing the information outside of the context of the guessing game; the upper part (a parallelogram) corresponds to the benefits that come purely from the guessing game.

units. Information on peer salary and peer seniority can be useful for career decisions such as salary negotiations, asking for a promotion or deciding whether to take an outside offer. However, learning about salaries may be subject to more substantive information frictions than learning about seniority. For example, there may be stronger privacy norms around salary than around seniority.

3.2 Training

To elicit valuations in an incentive-compatible way, we employ the Becker-DeGroot-Marschak (BDM) method. We use an open-ended variation (Andersen et al., 2006), in which the respondent bids against the computer for a particular item (e.g., a piece of information). The rules are as follows. The respondent's bid is compared to a price that is determined by a random number generator. If the respondent's bid is lower than the price, then the respondent gets a dollar amount equal to the price. If the bid is higher than the price, then the subject gets the item and no dollar amount. The rules of this mechanism thus encourage a dominant strategy in which respondents bid exactly their true valuation for the item. The rationale for this dominant strategy is equivalent to that in the Vickrey auction, wherein the dominant strategy is also to bid one's true valuation.

One important detail of the BDM mechanism is that, as explained in this training, all subjects must provide a bid for the item at hand, but this bid is not always "executed." We tell subjects that bids from "a few lucky participants" will be chosen at random to be executed. Subjects find out if their bids are selected on the screen immediately after entering their bids. For the "few lucky participants," the next screen also informs them about the outcome of the mechanism (i.e., whether they will receive the item or whether they will receive a sum of money to be deposited in their bank accounts). The survey then terminates prematurely, thereby excluding the participant from the subject pool. For those who are not among the "few lucky participants," the following screen notifies them that their bids will remain hypothetical. These subjects continue with the rest of the survey.

We do not specify to the respondent the number of participants whose bids are selected to be executed. In practice, we select 1% of the subjects invited to the survey. We select this small fraction for two reasons. First, the selected respondents cannot continue with the rest of the survey, so a higher share of respondents selected reduces the sample size. Second and most important, the firm wanted to limit the number of items being allocated, because some of these items could be distractive to the employees (e.g., revealing the employee's salary to five peers).

Another important feature of the BDM mechanism is that subjects never "lose" money, because they choose between receiving money or an item. Many studies use this type of

mechanism (Allcott and Kessler, 2018; Fuster et al., 2018), which differs from another common mechanism in which subjects must pay for the information out of their pockets. We do not implement this latter mechanism, because the firm wanted to avoid collecting payments from its employees' bank accounts.

Because subjects are probably not familiar with this method for eliciting willingness to pay or accept, we include instructions about this mechanism at the beginning of the survey. In our instructions, we note explicitly that it is in the respondents' best interest to bid their true valuations. Additionally, we include a couple of practice questions to familiarize subjects with this methodology.[11]

3.3 Guessing Game

Respondents must guess the average salary/seniority among a group of peers. Before providing their guesses, respondents receive a precise definition of the peer group. For example, a particular respondent was told that her peers were the "tellers from branch 130." To eliminate any remaining ambiguity, we also provide respondents with the total number of members in the peer group. The median size of the peer groups in our sample is 21 members, although the size varies widely from 6 to 89 members. To make the difficulty of the guessing game more comparable across peer groups, we asked each subject to guess the average salary or seniority among a set of five peers (excluding the respondent). The instructions of the guessing game include the list of first and last names of the five peers who were chosen randomly.

The guessing game offers a reward for accuracy: if the guess falls within 5% of the true average characteristic of the five peers, the subject receives an extra $\$X$ in payment from the experimenter, in addition to the other survey rewards. Because all employees had accounts at their employer's bank, we deposited the rewards automatically in the respondent's bank account within a couple of days. We randomized the size of the reward to take values $X \in \{\$13, \$26, \$39, \$52, \$65\}$, with equal probability. By randomizing the size of the reward, we generate exogenous variation in the expected benefits of holding accurate beliefs.

We give subjects three minutes to read the instructions and provide a guess. A clock in the upper left corner of the screen displays the time remaining. If the respondent does not provide a guess within the allotted time, the guess does not qualify for the reward. As we explain in detail in the following section, a key measurement of our study is the willingness

[11] One question elicits the willingness to pay for an iPhone X. The other question elicits the willingness to forfeit a lottery that pays \$100 with probability \$50. There are no correct or incorrect answers for these questions. If the subjects provide an abnormally high or abnormally low bid, we report back to them in the next screen, describing the typical responses to these questions .

to search for information when provided with additional time.

3.4 Confidence in the Initial Guess

To elicit the subjects' confidence in their belief, we elicit the probability with which they expect to win the reward. The subject can respond with any number from 0% to 100%, in 1% increments. Although this measure is not incentive-compatible, respondents have no obvious incentive to under- or over-report their true perceptions.

Nevertheless, to address this potential concern, we implement an alternative measure that is designed to be incentive-compatible. We elicit the willingness to forfeit the guessing game (i.e., the certainty equivalent that would make the subject forfeit the right to play the guessing game). Subjects can bid any dollar value from 0 to the prize amount. We use the BDM mechanism to incentivize truth-telling and explain to the subjects that it is in their best interest to bid their true willingness to pay for the guessing game. As previously explained, in our BDM elicitation, all individuals must provide bids, but the bids are executed only for a small sample selected at random for whom the survey is prematurely terminated.

We then normalize the bid that the subject provides as a share of the prize amount. This normalization makes the outcome more comparable across subjects who are randomly assigned to different reward amounts. This outcome takes a value from 0 to 1, like a probability. Indeed, if the subject were risk-neutral, this outcome should reveal the perceived probability of winning the guessing game.

Despite the advantage of being incentive compatible, this alternative measure has some disadvantages. First, it introduces measurement error to the extent that the willingness to forfeit the game depends not only on the perceived probability of winning but also on risk preferences. Second, it introduces measurement error in that bidding to have the right to play the game may be difficult to understand. We try to mitigate this problem by breaking the question into parts. First, we elicit the probability that respondents win the game first. Second, we calculate and show the subjects their expected value of the game (i.e., the subjective probability of winning the game multiplied by the reward amount). Last, we ask respondents to bid for the right to play the game.

3.5 Willingness to Search for Information

We want to elicit whether the respondents expect to be able to search for more information if given additional time. For instance, subjects may look up information on the Internet, ask one of the five peers in the list about their salaries, or ask information from other peers, managers, or human resources employees.

Subjects are told that some participants, selected at random, will be given the opportunity to get an extra week to search for information and revise their guesses. We ask respondents to report the likelihood that they could guess accurately if given the extra week. The difference between the winning probabilities with and without the extra week measures the expected value from searching (i.e., if subjects expect the probability of winning the guessing game to increase with the extra week, that would indicate that they expect to find useful information).

In addition to the non-incentivized version of this belief, we devise an incentivized (but more complex) version. We elicit the willingness to forfeit the guessing game conditional on getting the extra week. As usual, these bids are executed only for a minority of randomly selected respondents. The alternative measure of willingness to search for information is equal to the difference between the willingness to forfeit the guessing game with and without the extra week, again normalized as a share of the game reward.

3.6 Willingness to Pay for Readily Available Information

After subjects enter their guesses, they are given the chance to buy a piece of information, that is, a signal related to the object about which they are guessing. This signal consists of the average salary or seniority among different samples of five peers. Although imperfect, this piece of information is still useful to improve the accuracy in the guessing game and to learn about the average salary among all peers. To elicit this information in an incentive-compatible way, we employ the previously described BDM methodology: subjects enter their bids and compete with a bid generated by the computer. As usual, these bids are executed only for a minority of randomly selected respondents.

3.7 Willingness to Share Information with Peers

So far, we have explored the costs of finding out others' salaries. Now, we turn to the question of whether individuals face costs when they have to share information about their own salaries with others.

We elicit an objective and incentive-compatible measure of the cost, or benefit, of sharing information about one's own salary or seniority with peers. We inform respondents that we are considering sending an email to the five peers whose salaries or seniority they guessed. This email includes the first and last name of the respondent and the respondent's own salary or seniority. This email explicitly states that the information is being shared in the context of an experiment. Because the value of sending this email can be positive or negative, we elicit preferences about this email in two steps. First, we ask respondents whether they would like us to send this email. For respondents who want us to send it, we ask them to report their

willingness to pay for sending the email. For respondents who do not want us to send it, we ask them to report their willingness to accept payment in exchange for sending the email.[12] Because we wanted to make it clear to respondents that it was entirely up to them whether we send this email or not, we capped the range of the bids: the instructions noted that, by bidding $125, the respondent would get their wish (either to send or conceal the email) for sure. The resulting measure of willingness to share information can takes values from -$125 to $125. A positive amount indicates that the subject is willing to pay that amount to conceal her salary. A negative amount indicates that the subject is willing to pay (the absolute value of) that amount to reveal her salary.

3.8 Perceived Privacy Norms

The last screen of the survey shows three subjective questions related to privacy norms. The first question, *Unacceptable*, elicits the norm directly, asking whether it is "socially acceptable to ask someone about their salary or seniority", with possible answers of "highly acceptable," "somewhat acceptable," "somewhat unacceptable", and "highly unacceptable." One potential challenge with this measure is that an individual may perceive a certain norm and still feel comfortable breaking it. Thus, we include a second question, *Uncomfortable*, to elicit whether the respondent finds it "uncomfortable to ask information about salary/seniority to your peers" with the possible answers "not at all," "a little uncomfortable," "uncomfortable" and "very uncomfortable."

Last, individuals may be averse to asking about salary and seniority not because they want to avoid bothering others, but because they want to avoid being asked to reciprocate by revealing their own information. To assess this possibility, we ask the question *Reciprocal*: "if you ask a peer about his or her salary/seniority, would you expect this peer to ask you about your salary/seniority?" The possible answers are "Yes" and "No."

4 Institutional Context, Data, and Subject Pool

Our study was conducted in collaboration with a large, private, commercial bank in Asia with thousands of employees, hundreds of branches, and billions of dollars in assets and revenues. The setting has several features in common with firms of similar size around the world. For example, regarding the degree of pay inequality, the ratio between the 10th and 90th percentiles of salaries is 0.21 in this firm and, by comparison, is 0.19 for the average medium-sized U.S. firm (Song et al., 2016).

[12] As before, bids are executed only for a random minority of respondents and remain hypothetical for the rest.

The firm is typical in other relevant respects. It does not have an open salary policy, and its standard labor contract explicitly prohibits employees from sharing salary information. Many organizations around the world have similar policies, particularly in the United States (PayScale, 2018; Hegewisch et al., 2011). 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). Several other surveys corroborate this pattern of pay secrecy. A survey of about 1,000 companies indicates that only 3% have open salary policies (Scott, 2003), and less than a quarter disclose data on salary ranges. Moreover, a 2001 survey of U.S. employees finds that more than one-third are forbidden from discussing their pay with coworkers (Day, 2007; Vault, 2001).

These pay secrecy policies seem to be consequential. In our firm, survey data indicate that almost half of the employees never discuss their salaries with coworkers, and most employees prefer that the firm disclose average salaries by position titles (Cullen and Perez-Truglia, 2018). Other firms and countries report similar patterns. For example, according to a 2017 survey of Americans aged 18-36 years, 70% report that they never discuss their salaries with coworkers (WSJ, 2017). In a survey of employees from eight developed countries, most report that they are uninformed about salaries and want employers to be more transparent regarding pay (Glassdoor, 2016).

Last, in the firm, social norms seem to prevent employees from discussing pay with each other. According to our own survey data, 69% of employees believe that it is socially unacceptable to ask peers about their salaries. This secrecy norm is believed to be present in the United States (Edwards, 2005) and other countries. For example, a survey of 1,022 participants in the Northampton, UK, labor market found that less than half (48%) discuss salaries with their peers (Burchell and Yagil, 1997). Fox and Leshem (2005) presents survey data from Israel that most individuals report feeling highly uncomfortable when asked about their pay and other financial matters. Bierman and Gely (2004) describes the salary taboo in Canada. There is no consensus on the origins of this behavior or why it is so widespread around the globe. One argument is that pay can act as a signal of self-worth, and thus individuals feel that discussing their pay is equivalent to discussing their self-worth (Trachtman, 1999).

4.1 Survey Implementation

We start with the universe of thousands of employees. We focus on two specific units of the firm, with some added filters. For instance, to avoid any contamination, we exclude employees who participated in a previous survey that was related to peers' salaries (Cullen and Perez-Truglia, 2018). After filtering, we invited the remaining 1,899 employees to take the survey.

Appendix B includes a sample of the invitation email (stripped of formatting and identifying information). The survey was not compulsory, but employees were encouraged to participate. Indeed, the unit heads reached out to their employees by email to encourage participation in our survey. The invitation email did not provide any specifics about the content of the survey, but it explained that survey participants could earn monetary rewards, which would be deposited in their bank accounts, for participating in the survey.

The email invitations were sent gradually from February 9, 2018, to March 1, 2018. We sent a reminder by email to the subjects who had not completed the survey after one week of sending the original email and another reminder two weeks after the original email. The first subject responded on February 9, 2018, and the last subject responded on March 14, 2018. Of the 1,899 invitations sent, 752 individuals finished the survey, corresponding to a 39.6% response rate.[13] The median respondent took 16 minutes to complete the survey.

4.2 Descriptive Statistics and Randomization Balance

The subject pool includes employees from 45 different positions, such as tellers, salespeople, and branch directors. Of these, 18% are located in the two headquarter offices, and the rest are scattered across several branches.

Table 1 presents some descriptive statistics about the subject pool. Column (1) corresponds to the entire sample of 752 survey respondents: 72% are female, 86% finished college or a higher degree, and on average they are 29 years old and have been working at the firm for the last 4.2 years. In Appendix A.1, we show that this subject pool is representative of the universe of employees in these same observable characteristics.

Regarding pay inequality, the mean absolute difference between one's own salary and the average salary among all peers is 14% of one's own salary. In comparison, seniority has more horizontal inequality: the mean absolute difference between one's own seniority and the average seniority among all peers is 137% of one's own seniority.

We cross-randomized two features of the survey. In columns (2) and (3) of Table 1, we break down the descriptive statistics by the two survey types, salary and seniority. The last column reports p-values for the null hypothesis that the average characteristics are the same across these two treatment groups. The results show that, consistent with successful random assignment, the observable characteristics are balanced across the two treatment groups. The second feature of the survey that we randomized was the reward amount for the guessing game, which takes one of five different values. Columns (5) through (9) of Table 1 provide the corresponding balance test for this treatment arm. Again, the results are consistent with

[13] By construction, this sample excludes individuals who were randomly selected to have their surveys terminated prematurely (e.g., the subjects whose bids were selected to be executed).

successful random assignment.

5 Results

5.1 Misperceptions

Figure 2 shows misperceptions about average peer salary, comparing two different benchmarks in panels (a) and (b). This figure indicates that only 27% of subjects guess within 5% of the correct answer. The mean absolute difference between the perceived average and the actual average (i.e., the mean absolute error) is 16%. These misperceptions are economically significant, but they are not skewed: approximately as many people overestimate the average peer salary as the number of people who underestimate it, resulting in an average underestimation of peer salary of just 0.06%, which is statistically indistinguishable from zero (p-value=0.490).

Only 24% of respondents guess an average peer salary within 5% of their own salary. This evidence suggests that they use other information sources besides their own salaries. To assess whether this extra information improves their accuracy, Figure 2.a compares the misperceptions with respect to the benchmark scenario in which individuals report their own salaries as their guesses. The extra information does not seem to improve their accuracy: according to a non-parametric test reported in Figure 2.a, we cannot reject the null hypothesis that these two distributions of misperceptions are the same (p-value=0.071). If individuals report their own salaries as their guesses for average peer salary, the mean absolute error (14%) would be, if anything, slightly lower than the mean absolute error of the actual guesses provided by the subjects (16%).

Figure 2.b provides another useful benchmark: what misperceptions would look like if an individual's guess for average salary among the five peers equals the actual average salary among all peers. The MAE would have been much lower (7%, instead of 16% in reality). This finding shows that most misperceptions are not caused by asking about a specific subsample of peers.

In Cullen and Perez-Truglia (2018), we collect some measurements that are related to the ones presented in this section, such as the degree of misperceptions. Any comparisons of the results across the two surveys must account for differences in samples, definitions, and methods. Despite these caveats, however, the results from this survey are roughly consistent in magnitude with the corresponding results from Cullen and Perez-Truglia (2018).[14]

[14]For example, in Cullen and Perez-Truglia (2018), we elicit beliefs about the average salary among all peers. In this survey, we elicit beliefs about the salary of a specific set of five peers, which is by construction a harder object to guess. Consistent with this interpretation, the MAE of peer average guesses are higher in

5.2 Perceived Accuracy

Figure 3.a shows the distribution of the self-reported probability of winning the guessing game with the initial guess. Employees understand that they do not have perfect beliefs, but many (56%) believe that their guess for average peer salary is within 5% of the truth. Yet, only 27% of guesses actually fall within that range, indicating overconfidence among respondents.

It is possible that individuals misreport their perceived accuracy. For example, some subjects may over-report their true confidence to impress the surveyor. As a robustness check, we present results with our incentive-compatible proxy: willingness to forfeit the guessing game as a share of the game reward. If the subject is risk-neutral, this proxy equals the expected probability of winning the game. Indeed, this incentive-compatible proxy is positively and significantly correlated to the self-reported measure (correlation coefficient of 0.20, with a p-value<0.01).[15] The distribution of this incentive-compatible proxy for the probability of winning the game, shown in Figure 3.b, is roughly comparable to the distribution of the self-reported equivalent, shown in Figure 3.a. Last, we still find that individuals are overconfident about their own accuracy if we use the incentive-compatible proxy instead of the self-reported measure.[16]

5.3 Willingness to Search for Information in the Wild

In addition to the probability of winning the game without the extra week, Figures 3.a and 3.b show the expected probability of winning the game with the extra week. Regardless of whether we use the self-reported (Figure 3.a) or incentive-compatible (Figure 3.b) measures, employees tend to expect their accuracy to increase substantially with the additional week. Indeed, according to Figure 3.a, the average probability of winning the game increases from 56% to 79% (p-value of the difference is <0.001) with the additional week. Although somewhat smaller in magnitude, this gap remains significant when using the incentive compatible measure: the average probability increases from 43% to 57% (p-value<0.001).

Figure 4.a shows one standardized way of measuring the willingness to search. This measure equals the difference between the probability of winning the game with and without the extra week, divided by the probability of losing the game without the extra week. This

this survey (16%), relative to the MAE reported in Cullen and Perez-Truglia (2018) (12%).

[15] We do not expect the correlation to be perfect. First, the proxy equals the self-reported probability only if the individual is risk-neutral, but in practice, different individuals may have different risk aversions. Second, there is probably significant measurement error in both of these variables, particularly the incentive-compatible one, as it is more complex to understand and thus more prone to errors.

[16] On average, 43% of subjects expect their guess for average peer salary to be within 5% of the truth, which is still significantly higher than the 27% of guesses that actually fall within that range.

sample excludes 30 individuals who reported 100% confidence in their initial guesses and for whom the denominator would be zero. A 0% in this measure means that the individual does not expect to eliminate any of the initial inaccuracy, and 100% means that the individual expects to fully eliminate the initial inaccuracy.

There is large variation in this measure of willingness to search. The sample can be roughly divided in three thirds. The first third does not expect to get better (i.e., less than 10% better) with an extra week. For those individuals, the marginal costs of searching for information must be greater than the marginal benefits from winning our guessing game. The second third of the sample expects to improve all the way up to certainty. The misperceptions for these individuals are largely voluntary (i.e., they do not acquire information, because the benefits from the information are not significant enough yet). The last third of the sample is between these two extremes.

It must be noted, however, that the anticipated gains may not coincide with the actual ex-post gains from searching. In particular, given that individuals are overconfident about the accuracy of their initial guesses, they also may be overconfident about the expected gains from searching.[17] In any case, our measure of anticipated gains is the relevant measure for the decision to search for information or not.

Our measure of willingness to search may overestimate the "natural" willingness to search for salary information, because our game may provide some individuals with an "excuse" to ask peers about their salaries. For example, participants could motivate their request for information by mentioning that they want to win the guessing game, which may be a more acceptable justification than ordinary alternatives.

This evidence that some individuals are willing to search for information when provided with benefits is consistent with rational inattention models, according to which individuals wait to acquire information until the benefits of being informed are high enough. Moreover, the randomization of the rewards provides a test of the rational inattention model: a higher reward should cause individuals to search more intensively. Figure 4.b illustrates this test. Consistent with rational inattention, individuals who are assigned to higher rewards expect to search more intensively for information. More precisely, a $100 increase in the guessing reward results in an expected accuracy increase of 15 percentage points, a substantial expected gain that is equivalent to 27% of the average expected accuracy rate of 56 percentage points.

[17] We could have also measured the ex-post gains from searching by giving the extra week to a significant fraction of the sample. However, we refrained from doing this due to institutional constraints.

5.4 Willingness to Pay for Readily Available Information

The previous section indicates that individuals are more willing to search when provided with incentives, but it does not address the gross benefits of the sought-after information. As described in Section 2, these benefits can be assessed by means of the willingness to pay for information.

Figure 5.a shows the distribution of willingness to pay for a signal indicating the average salary among a different sample of five peers. The median employee is willing to pay roughly \$13 for the information. This amount suggests that the bottom half of the subjects, who are willing to pay no more than \$13 for the information, have misperceptions mostly due to lack of interest.

The upper part of the distribution, however, is willing to pay substantial amounts. For example, the top 26% of subjects are willing to pay amounts that are approximately uniformly distributed between \$100 and \$1,300, with a median of \$652 and a mean of \$640. These large valuations suggest that these employees have misperceptions not because they do not care about the information but because the costs of searching for the information are high. These subjects may need to use the information for high-stakes decisions, such as whether to take an outside job offer or request a raise or a promotion. For instance, suppose that an employee wants to buy information about peer salary for use in salary negotiations. If the information is expected to translate into a salary increase of just 5% and for just one year, then the employee should be willing to pay more than two weeks' worth of salary for this information (Stigler, 1962).

For individuals in the upper part of the distribution, the willingness to pay for information is an order of magnitude higher than the guessing game rewards. These individuals are not bidding for the information with the main goal of winning the guessing game. However, it is plausible subjects bidding close to the median (\$13) are bidding primarily to improve the chances of winning the guessing game. We can test the rational inattention hypothesis, according to which individuals should be willing to bid more for information when they stand to gain more from it. Figure 5.d reports the results for this test. We find that, consistent with rational inattention, increasing the reward size by \$1 increases the median willingness to pay for information by \$0.38 (p-value=0.029).

One potential concern is that our estimates of willingness to pay may be sensitive to the elicitation method. Evidence from a previous study (Cullen and Perez-Truglia, 2018) suggests that this may not be a source for concern. When measuring willingness to pay for information about peer salary using the price-list method, instead of the open method used in this experiment, we find a willingness to pay for information in the same order of magnitude. This finding is consistent with prior evidence that measures of willingness to pay are similar

across different elicitation methods (Brebner and Sonnemans, 2018).

5.5 Willingness to Share Information with Peers

This evidence suggests that individuals face significant costs when acquiring information from others. In turn, this section discusses whether they also face frictions when sharing information with others.

Figure 6.a shows the distribution of the willingness to share own salary information with the sample of five peers. There is large variation in this outcome. On the one hand, roughly 20% of employees are willing to pay the experimenter to send an email revealing their salaries to peers, a surprising result considering that individuals could in principle reveal the information themselves for free (e.g., by sending a similar email or by mentioning the information in casual conversations). However, we explicitly inform respondents that the email revealing their salaries will mention that it is being sent in the context of an experiment. Thus, individuals may be willing to pay for this email so that they can use the experiment as an "excuse" to reveal their salaries. This evidence suggests that our measure of willingness to share information may overestimate the "natural" willingness, because our game may provide a novel "excuse" to share the information.

On the other hand, roughly 80% of respondents are willing to pay to avoid sending the email. Moreover, even within each of these two groups, there is large variation in preference strength. Roughly 40% of subjects are mostly indifferent about their privacy, in the sense that they are willing to pay less than $5 to reveal or conceal their salary information. The remaining 60%, however, show strong preferences: 4% of respondents are willing to pay more than $125 to reveal their information, but a whopping 40% required more than $125 to allow the experimenter to reveal their salaries.

In addition to being averse to sharing information about their own salaries, employees also may be averse to sharing other types of salary information. This aversion could explain the experimental evidence from Cullen and Perez-Truglia (2018) that employees who receive information about peers' or managers' salaries do not share that information with peers, not even with their closest peers in the network.

This evidence suggests a privacy norm for salary information. This interpretation is consistent with the subjective data, presented in Figure 7. Figure 7.a shows that 69% of employees find it unacceptable to ask coworkers about their salaries. The results are similar under the alternative subjective measure: Figure 7.b shows that 53% of employees find it uncomfortable to ask coworkers about their salaries.

This variation in willingness to share one's own salary with coworkers may reflect a direct preference for privacy, but it also may respond to strategic incentives (i.e., individuals may

expect positive or negative effects from revealing their salaries to coworkers). For example, if an employee reveals to a coworker that she gets paid more, her peers may stop treating her well, or if her manager finds out, the manager may deny her a raise. According to this rationale, the higher the relative salary of the employee within the peer group, the stronger the preference should be to keep the salary information private. Yet, the status models (Frank, 1984; Bursztyn et al., 2018) predict the opposite: employees with higher relative salary should be more excited about revealing their salary, because that can result in higher status.

Our unique data on willingness to pay for privacy allows a direct comparison of these two mechanisms. Figure 6.b shows the relationship between the willingness to pay for privacy and the perceived distance between own-salary and the reference group. There is a significant relationship: increasing the individual's perceived relative salary by 1 standard deviation is associated with an increase in willingness to pay for privacy of \$8.4, which is equivalent to 11% of the standard deviation of this outcome (\$75), while the relationship is exactly reversed, though less precisely estimated, for those asked about seniority. A rise in perceived relative seniority of 1 standard deviation is associated with a decrease in the demand for privacy of \$8.1, and those who believe they are more senior than the average of the reference group have a positive willingness to pay for us to send an email about their tenure. This evidence suggests that status concerns are likely at play in determining who chooses to announce their seniority, but not in determining who reveals their own salary. A perception of earning more is a deterrent for sharing salary information, consistent with the notion that this information could have detrimental effects on team effort (Cullen and Perez-Truglia (2018), cause resentment or put at risk their exceptionally high pay.

The fact that a large fraction of individuals strongly prefer not to share their salary with others suggests a demand for privacy. This demand for privacy may be an important contributing factor in the lack of information diffusion. Consistent with this interpretation, survey data indicate that employees are more comfortable with sharing salary information when they can do so anonymously. For instance, surveys from seven developed countries indicate that more than 62% of employees would be willing to share information about their own salaries if they could do so anonymously (Glassdoor, 2016). Preferences for anonymous sharing suggest that the desire to avoid personal resentment may partially drive the demand for privacy among higher-paid employees. Anonymity does not necessarily offer a solution to team morale concerns or protect information that can be used to help others renegotiate.

We also collect subjective data on the expected reciprocity. About 89% of employees believe that asking others about their salaries will result in others asking for the same information. This evidence suggests an additional reason why privacy preferences may dissuade

individuals from searching for salary information: individuals who want to avoid being asked about their own salaries will avoid asking coworkers about their salaries.

Importantly, sending messages about salary on behalf of participants allows us to remove complications created by a norm of reciprocation. Higher-paid employees may be motivated by status concerns to reveal, but nevertheless fear embarrassing lower-paid peers who reciprocate disclosure. However, our results show that higher-paid employees are willing to pay for privacy at higher rates relative to peers even when they can disclose their salary information via email without the expectation of reciprocation. This is stronger evidence that status concerns associated with salary disclosure are muted than we could get through incentivizing direct in-person exchanges.

5.6 Diffusion of Information about Salary Versus Seniority

We start by comparing the misperceptions between average peer salary and average peer seniority. Within-group variation is higher with seniority than with salary, so guessing the average seniority of five peers is significantly harder than guessing the average salary of five peers. To make the degree of misperceptions comparable, we first standardize them. We define the standardized absolute error as the absolute difference between the guess for average salary (seniority) and the actual average salary (seniority) of the five peers, divided by the within-group standard deviation in salary (seniority).

Table 2 presents regressions of some key outcomes on a dummy variable that equals 1 if the survey type is seniority and 0 if the survey type is salary. Columns (1) and (2) measure differences in misperceptions. To validate the standardization, column (1) uses as a dependent variable the hypothetical standardized absolute error based on the hypothetical scenario in which respondents report their own salary or seniority as their guess. There is no significant difference in these hypothetical misperceptions: the standardized absolute error is 0.724 for salary and 0.704 for seniority, with a p-value of the difference of 0.662. In other words, if they only have access to their own salary or seniority information, subjects are equally accurate in guessing their salary or seniority. This finding confirms that the standardized measure of misperceptions provides a fair comparison between the two types of information.

In turn, the dependent variable in column (3) of Table 2 is the standardized absolute error of their actual guesses. According to this measure of misperceptions, subjects are much better at guessing the average seniority than they are at guessing the average salary. The mean absolute error equals 1.069 standard deviations for salary but just 0.525 standard deviations for seniority (difference p-value < 0.001). The comparison between columns (2) and (3) suggest that, when guessing salary, subjects do not do better than if they only had information about their own salaries, and when guessing seniority, subjects do substantially

better than they would if they only had access to information about their own seniority. This evidence suggests that there is substantially more information diffusion for seniority than for salary.

One possible interpretation for the differences in misperceptions is that information frictions are higher for salary than for seniority (i.e., the marginal cost curve for salary is higher than the marginal cost for seniority). Alternatively, this difference could indicate that the marginal benefit curve is higher for seniority than for salary. To distinguish between these two explanations, we exploit data on the willingness to pay for readily available information. If the differences in misperceptions are driven by higher marginal benefits for seniority, we should observe a higher willingness to pay for information about seniority than about salary. Column (3) of Table 2 tests that hypothesis. We find that, on average, there is a large difference, but in the opposite direction: employees are willing to pay more for a signal of peer salary ($179) than for a signal of peer seniority ($130), with the difference being statistically significant (p-value=0.030). This evidence suggests that the differences in misperceptions between salary and seniority must be due, at least partially, to differences in the marginal costs (i.e., information frictions).

Several reasons may explain the higher frictions for learning about salaries, compared to seniority. First, the firm may discourage employees from discussing salaries but not seniority. We believe that this scenario is unlikely, because we tacitly lifted the non-disclosure policy for the duration of our study. Also, about half of the employees report that they do discuss salaries with their coworkers, although they have nothing to gain from revealing that. This evidence suggests that the firm's disclosure policy may not be the main concern in the employee's mind.

Our preferred explanation is that employees find it easier to learn about seniority than it is to learn about salaries because of differences in privacy norms. We present two pieces of suggestive evidence on this regard. We start with the subjective data, presented in Figure 7. These data suggest that salary is a much more sensitive topic than seniority: 69% of employees find it unacceptable to ask a coworker about salary, compared to only 6% who find it unacceptable to ask about seniority. Similarly, 53% of employees find it uncomfortable to ask a coworker's salary, whereas only 5% find it uncomfortable to ask about seniority. The difference in these distributions across salary and seniority are highly statistically significant, according to the non-parametric tests reported in figure 4 (p-value<0.001 for both *Unacceptable* and *Uncomfortable*). On a scale from 0 to 3, *Unacceptable* averages 1.8 for salary versus 0.5 for seniority, and *Uncomfortable* averages 1.6 for salary versus 0.3 for seniority.[18]

[18]Norms about reciprocity are similar: 89% of respondents report that they will get asked to reveal their own salaries if they ask someone about theirs, whereas 93% of respondents report that they will be asked to reciprocate when asking someone about their seniority.

The second piece of suggestive evidence is based on the revealed-preference measure on willingness to pay for privacy, reported in column (1) of Table 2. There is an economically large and statistically significant (p-value<0.001) difference in this outcome. Individuals need compensation of $67, on average, to allow an email revealing their salary to peers. Individuals need compensation of only $28, on average, to allow an email revealing their seniority. This large difference also suggests that the differences in misperceptions between salary and seniority are due in part to differences in privacy norms.

5.7 Gender Differences

In the case of salary information, there is a widespread belief that pay secrecy tends to hurt women and minorities disproportionately, in part because women are less likely to search for salary information (Babcock and Laschever, 2009). The only evidence on this matter, however, comes from survey data. For example, in the United States, 65% of men and 53% of women believe that they have a good understanding of how people are compensated at all levels in their company. This gap is qualitatively consistent in eight countries included in the survey (Glassdoor, 2016). In a survey conducted by Cullen and Pakzad-Hurson (2017), participants of both genders believe that men are more likely than women to ask about and discover a co-worker's wage.

Some evidence suggests gender differences in the diffusion of some forms of information. For example, (Beaman et al., 2018) find that diffusion of productivity-enhancing information does not extend far beyond the initial individuals contacted; thus, women who happen to be peripheral in this network are less informed than men. Similar evidence indicates that job referral networks that operate through word-of-mouth tend to favor men over women (Beaman et al., 2018). However, there is no data on gender differences when it comes to the diffusion of salary information.

Table 3 presents regressions of several outcomes on a dummy variable that equals 1 if the employee is female and 0 if the employee is male. The evidence from Table 3 suggests that, consistent with the aforementioned survey data from (Glassdoor, 2016) and Cullen and Pakzad-Hurson (2017), women tend to be less confident than men about their knowledge of peer salaries. Column (9) indicates that, according to the self-reported measure, the perceived probability of winning the game is 60% for men versus 54% for women (p-value of the difference = 0.180). Column (10) indicates that, according to the incentive-compatible measure, the perceived probability of winning the game is 51% for men versus 40% for women (p-value of the difference = 0.028). However, the comparison of actual accuracy indicates that these differences in perceived accuracy are misleading. Column (8) of Table 3 indicates that, if anything, women are more accurate than men when it comes to guessing salaries: the

share of men winning the guessing game is 21.5% versus 28.6% for women (p-value of the difference = 0.159).

Besides this difference in confidence, there are no significant gender differences in any of the other outcomes reported in Table 3. Female and male respondents find it similarly acceptable to ask peers about their salaries. They are equally comfortable asking peers about their salary and find it equally likely that they will be asked to reciprocate. They have similar mean errors and mean absolute errors in their perceptions, and they are equally willing to search for information in the wild, to pay for readily available information, and to disclose their salaries to peers.

6 Conclusions

Barriers to the diffusion salary information can have significant implications for a broad range of labor market phenomena. Despite these implications, there is little direct evidence on how salary information diffuses in practice. We designed and implemented a field experiment with employees of a real organization to address this question.

We show that individuals have significant misperceptions about their peers' salaries. We show that some of these misperceptions are due in part to high search costs. Moreover, we provide suggestive evidence that the salary taboo at least partially drives these search costs. For example, we show that some individuals have a strong preference for concealing their own salary. We also show that these privacy norms, and the resulting misperceptions, are significantly lower when individuals have to learn about seniority instead of salary.

Several policies have been enacted to increase pay transparency. Our evidence suggests that some of these policies may benefit employees, helping them make informed choices and perhaps even gain leverage in salary negotiations. Moreover, our findings can inform the design these policies. For example, from 2016 to 2018, 13 of the 50 U.S. states passed legislation punishing employers that retaliate against workers who discuss wages with coworkers. Our evidence suggests that this policy alone may have a limited effect on how often employees discuss salaries, because they remain constrained by privacy norms. Indeed, in our own study, although we tacitly lifted the non-disclosure policy for the duration of our study, we find that a significant fraction of the sample was unwilling to search for and share salary information.

Disclosure policies have also been promoted with the intention of reducing wage discrimination against women. Indeed, information frictions played a central part in the Lilly Ledbetter Fair Pay Act of 2009. In this landmark case against pay discrimination, an anonymous note revealing the higher salaries of Lilly Ledbetter's peers proved pivotal in identifying in-

equity and retribution. The information was discovered nineteen years after Ledbetter joined the firm. The Fair Pay Act thus extended the statute of limitations for filing an equal-pay lawsuit, precisely because, as in the Ledbetter case, the discovery of unequal coworker salaries can take decades. More recently, several countries around the world, including the United Kingdom and Germany, have passed legislation mandating disclosure of salary information specifically designed reduce the gender wage gap (WGEA, 2017). Further evidence is needed, but our findings constitute suggestive evidence that disclosure policies may have a limited effect on the gender gap in salary information, as women and men seem to face similar information frictions. If anything, the only gender difference that we find is that women tend to be less confident than men about their knowledge of peer salaries.

Our evidence on employees' demand for privacy suggests that employees can be hurt significantly by transparency policies that result in disclosing non-anonymous information. Indeed, data from a separate survey of 2,033 employees from the same firm indicate that employees would choose higher transparency only if it is anonymous: Whereas 65% of respondents report that they would be better off if the bank disclosed average salaries by position, only 13% would be better off if the bank disclosed salary information in a non-anonymized manner. Disclosure policies sometimes allow for anonymity. For example, in 2018, California began requiring that employers provide prospective employees with their current employees' salary range (Pender, 2017; Siniscalco et al., 2017). However, many recent transparency policies involve non-anonymized information. For instance, in multiple U.S. states including California, Florida, and New York, employers must disclose the full names and salaries of all public employees on the Internet. Our findings suggest that these policies may need to be redesigned, for example, by removing the personally identifiable information (e.g., employees' names) from the public records.

References

Acquisti, A., C. Taylor, and L. Wagman (2016). The Economics of Privacy. *Journal of Economic Literature 54*(2), 442–492.

Adjerid, I., A. Acquisti, L. Brandimarte, and G. Loewenstein (2013). Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. *Proceedings of the Ninth Symposium on Usable Privacy and Security*, 9:1–9:11.

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

Allcott, H. and J. B. Kessler (2018). The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons. Working paper, National Bureau of Economic Research.

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

Athey, S., C. Catalini, and C. Tucker (2017). The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. *NBER Working Paper No. 23488*.

Babcock, L. and S. Laschever (2009). *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press.

Bartoš, V., M. Bauer, J. Chytilová, and F. Matějka (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review 106*(6), 1437–1475.

Bass, F. M. (1969). A New Product Growth for Model Consumer Durables. *Management Science 15*(5), 215–227.

Beaman, L., A. Dillon, and A. D. Lori Beaman (2018). Diffusion of agricultural information within social networks: Evidence on gender inequalities from Mali. *Journal of Development Economics 133*, 147–161.

Beaman, L., N. Keleher, and J. Magruder (2018). Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Journal of Labor Economics 36*(1), 121–157.

Bierman, L. and R. Gely (2004). Love, sex and politics? Sure. Salary? "No Way": Workplace social norms and the law. *Berkley Journal of Law and Employment 25*(1), 167–191.

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

Breza, E., S. Kaur, and Y. Shamdasani (2018). The Morale Effects of Pay Inequality. *The Quarterly Journal of Economics 133*(2), 611–663.

Brickley, J. A., C. W. Smith, and J. L. Zimmerman (2007). *Managerial economics and organizational architecture*. Boston: McGraw-Hill Irwin.

Burchell, B. and D. Yagil (1997). Socioeconomic and Political Initiators of Pay Comparison. *Work, Employment, & Society 11*(4), 737–748.

Bursztyn, L., B. Ferman, S. Fiorin, M. Kanz, and G. Rao (2018). Status Goods: Experimental Evidence from Platinum Credit Cards. *The Quarterly Journal of Economics 133*(3), 1561–1595.

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

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

Colella, A., R. Paetzold, A. Zardkoohi, and M. Wesson (2007). Exposing pay secrecy. *Academy of Management Review 32*, 55–71.

Corbett, W. R. (2002). Waiting for the Labor Law of the Twenty-First Century: Everything Old Is New Again. *Berkeley Journal of Employment and Labor Law 23*(2), 259–306.

Cullen, Z. and B. Pakzad-Hurson (2017). Equilibrium effects of pay transparency. *Working Paper*.

Cullen, Z. and R. Perez-Truglia (2018). How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*.

Danziger, L. and E. Katz (1997). Wage Secrecy as a Social Convention. *Economic Inquiry 35*(1), 59–69.

Day, N. E. (2007). An investigation into pay communication: Is ignorance bliss? *Personnel Review 36*, 739–762.

Edwards, M. A. (2005). The Law and Social Norms of Pay Secrecy. *Berkeley Journal of Employment and Labor Law 26*, 41–63.

Ellingsen, T. and b. Rosén (2003). Fixed or Flexible? Wage-Setting in Search Equilibrium. *Economica 70*(278), 233–250.

Ellison, G. and D. Fudenberg (1995). Word-of-Mouth Communication and Social Learning. *The Quarterly Journal of Economics 110*(1), 93–125.

Fox, S. and L. Leshem (2005). Displeasure From Disclosing Information on Financial Matter to Different Categories of Others. *Megamot 44*(1), 163–178.

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

Fuster, A., R. Perez-Truglia, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *NBER Working Paper No. 24767*.

Gely, R. and L. Bierman (2003). Pay secrecy/confidentiality rules and the National Labor Relations Act. *Journal of Labor and Employment Law 6*, 120–156.

Glassdoor (2016). Global Salary Transparency Survey. Technical report.

Goldfarb, A. and C. Tucker (2012). Shifts in Privacy Concerns. *American Economic Review 102*(3), 349–353.

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

Kuhn, P. and W. Gu (1998). A Theory of Holdouts in Wage Bargaining. *American Economic Review 88*(3), 428–449.

Kuhn, P. and W. Gu (1999). Learning in Sequential Wage Negotiations: Theory and Evidence. *Journal of Labor Economics 17*(1), 109–140.

Lawler, E. J. (2003). Pay Practices in Fortune-1,000 Corporations. *WorldatWork Journal 12*(4), 45–54.

Mankiw, N. G. and R. Reis (2002). Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve. *The Quarterly Journal of Economics 117*, 1295–1328.

Mas, A. (2016). Does Disclosure affect CEO Pay Setting? Evidence from the Passage of the 1934 Securities and Exchange Act. *Working Paper*.

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

Michelacci, C. and J. Suarez (2006). Incomplete Wage Posting. *Journal of Political Economy 114*(6), 1098–1123.

Mobius, M., T. Phan, and A. Szeidl (2015). Treasure Hunt: Social Learning in the Field. *NBER Working Paper No. 21014*.

Mobius, M. and T. Rosenblat (2014). Social Learning in Economics. *Annual Review of Economics 6*(1), 827–847.

Moellers, C., H.-T. Normann, and C. M. Snyder (2017). Communication in vertical markets: Experimental evidence. *International Journal of Industrial Organization 50*, 214–258.

PayScale (2018). 2018 Compensation Best Practices Report. Technical report.

Pender, K. (2017). New law bans California employers from asking applicants their prior salary. *https://www.sfgate.com/business/networth/article/New-law-bans-California-employers-from-asking-12274431.php*.

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

Phillips, M. (2009). Now Comes Lilly Ledbetter. *https://obamawhitehouse.archives.gov/blog/2009/01/25/now-comes-lilly-ledbetter*.

Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics 53*(8), 1761–1800.

Scott, D. (2003). Survey of Compensation Policies and Practices.

Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Siniscalco, G. R., E. M. Connell, and C. Smith (2017). State Pay Equity Laws: Where a few go, many may follow. *Mimeo*.

Song, J., D. J. Price, F. Guvenen, N. Bloom, T. von Wachter, F. Guvenen, D. J. Price, and J. Song (2016). Firming Up Inequality. *NBER Working Paper No. 21199*.

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

Trachtman, R. (1999). The Money Taboo: Its Effects in Everyday Life and in the Practice of Psychotherapy. *Clinical Social Work Journal 27*(3), 275–288.

Vault (2001). More employers ducking pay confidentiality issue. *Business and legal reports, Vault.com*.

WGEA (2017). International Gender Reporting Schemes. *Report of the Workplace Gender Equality Agency*.

Woodford, M. (2001). Imperfect Common Knowledge and the Effects of Monetary Policy. *NBER Working Paper No. 8673*.

WSJ (2017). Ask Me How Much Money I Make. *The Wall Street Journal, Oct. 26 2017.*.

Figure 1: Conceptual Framework

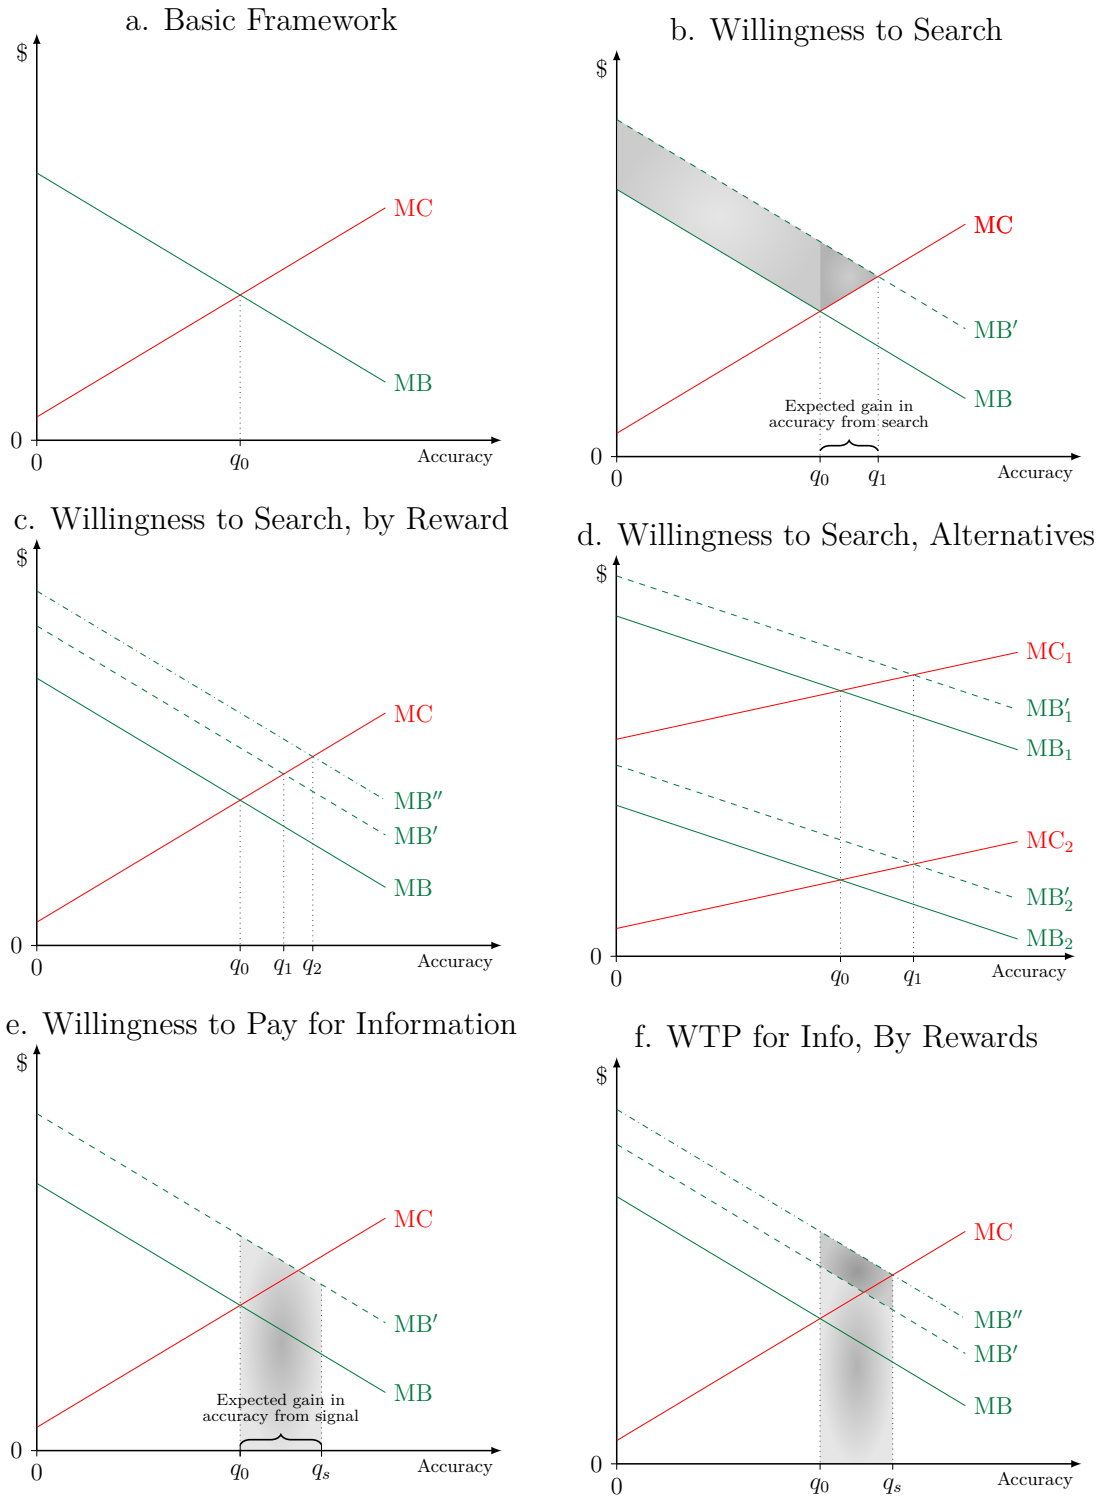

Notes: MC corresponds to the marginal cost of searching for information in the wild, while MB corresponds to the marginal benefits. The shifts from MB to MB′ represents the unexpected introduction of the guessing game. The shift to MB″ represents the introduction of a game with a higher reward amount.

Figure 2: Salary Information: Misperceptions

a. Guess Vs. Own-Salary Benchmark

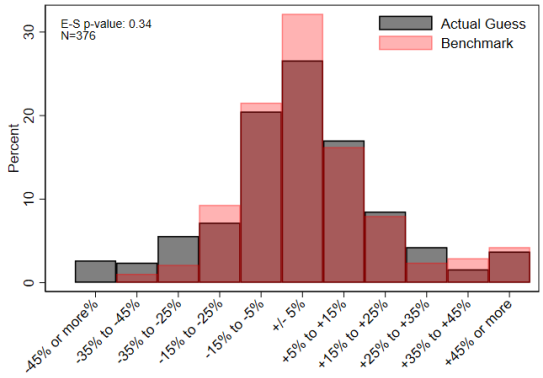

b. Guess Vs. Peer-Group Benchmark

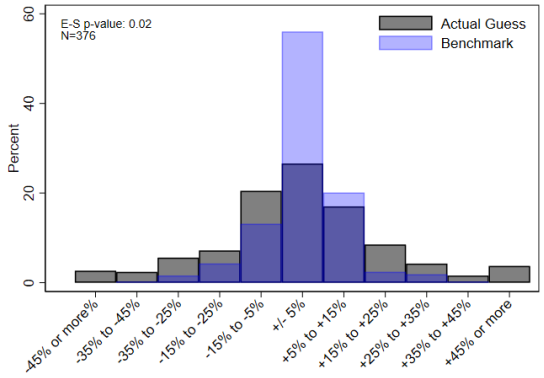

Notes: Histograms of the salary misperceptions, defined as the difference between the respondent's guess about the average salary among the sample of five peers (according the incentivized elicitation) and the actual average salary (according to the firm's administrative records), divided by the actual average salary. Panel (a) provides the following benchmark: what the misperceptions would have been if the respondent had provided a guess equal to her own salary. Panel (b) provides yet another benchmark: what the misperceptions would have been if the respondent had provided a guess equal to the actual average salary among all peers (not only the five selected peers).

Figure 3: Salary Information: Perceived Accuracy with and without the Extra Week

a. Self-Reported

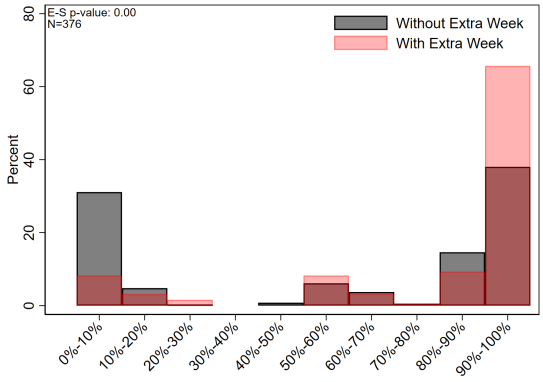

b. Incentive-Compatible

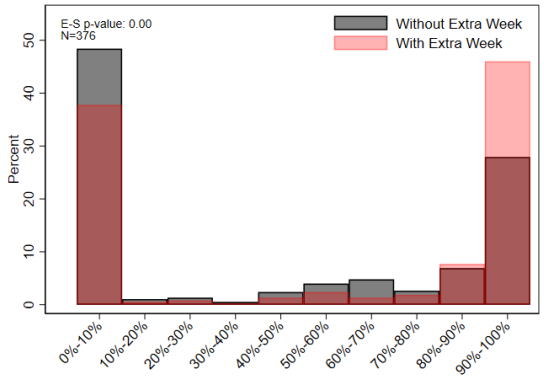

Notes: Panel (a) shows a histogram of the respondent's perceived probability of winning the guessing game, without an extra week to search for information (grey bars) and with the extra week (red bars). Panel (b) is equivalent to the first panel, only that instead of using self-reported probabilities, we use an incentive-compatible proxy: the ratio between the willingness to forfeit the guessing game and the reward amount.

Figure 4: Salary Information: Willingness to Search for Information in the Wild

a. Willingness to Search

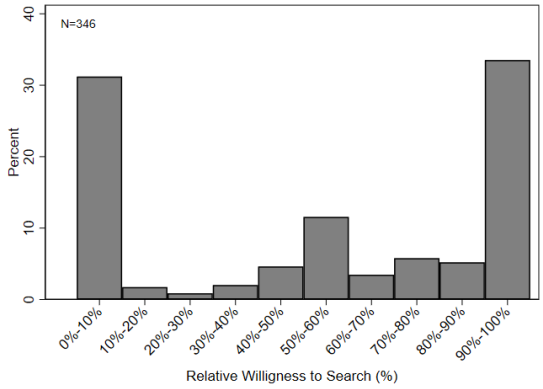

b. Rational Inattention Test

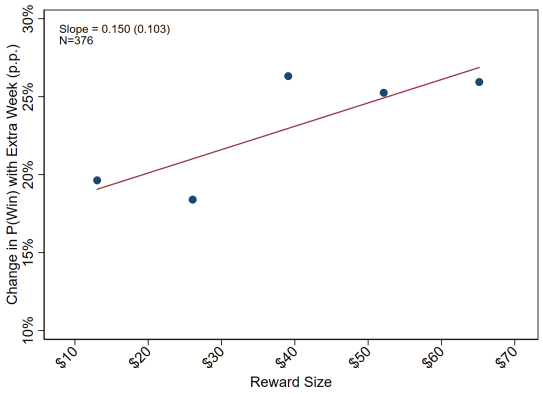

Notes: Panel (a) shows the difference between the probability of winning the game with and without the extra week, divided by one minus the probability without the extra week (this sample excludes 30 individuals who were 100% confident in their initial guess). Panel (b) provides a binned scatterplot with the relationship between the reward amount (x-axis) and the expected accuracy gain with the extra week (in percentage points). The slope is calculated with ordinary least squares, with standard errors clustered at the position level reported in parentheses.

Figure 5: Salary Information: Willingness to Pay for Readily-Available Information

a. Willingness to Pay for Information

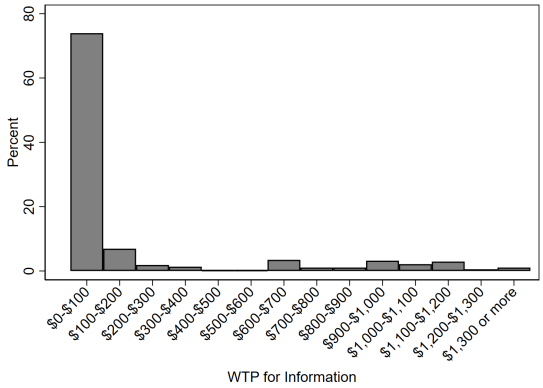

b. Rational Inattention Test

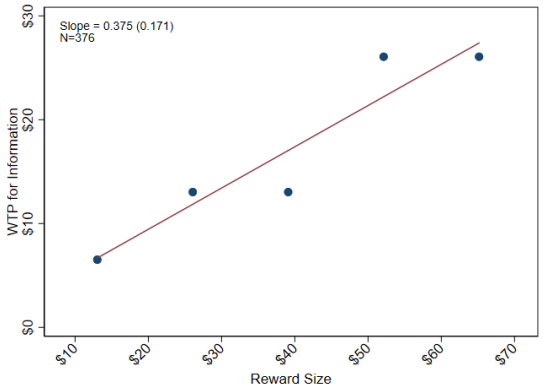

Notes: Panel (a) shows the distribution of the willingness to pay for information about the average salary among a sample of five peers, as measured by the respondent's incentive-compatible bid. Panel (b) provides a binned scatterplot with the relationship between the reward amount and the median willingness to pay for information. The slope is calculated with a quantile regression, with standard errors clustered at the position level reported in parentheses.

Figure 6: Salary Information: Willingness to Share Information with Peers

a. Willingness to Share Information

b. Association to Relative Salary

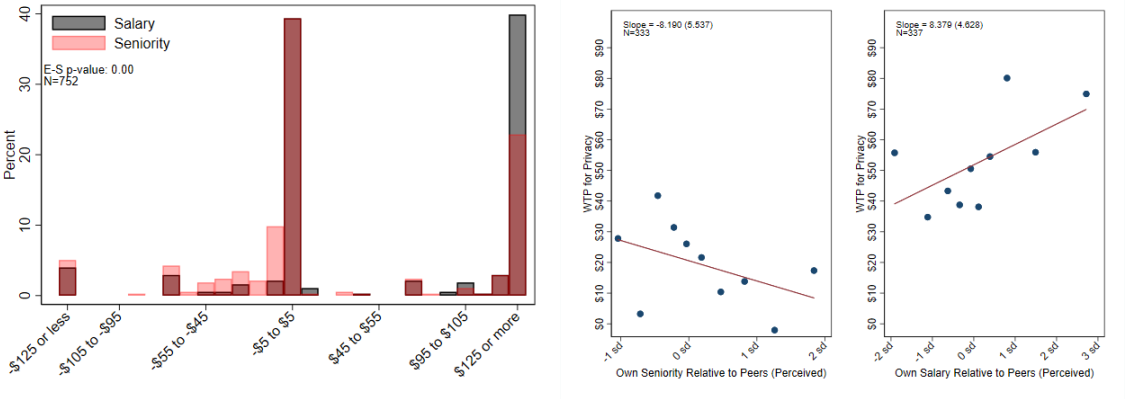

Notes: Panel (a) shows the distribution of the willingness to pay for privacy: negative values denote the amount the individual is willing to pay to reveal her information to peers, while positive values denote the compensation the individual is willing to give up in order to conceal this information. Panel (b) provides a binned scatterplot with the relationship between the willingness to pay for privacy and the respondent's perceived relative salary (or tenure) with respect to the reference peer group. The slope is calculated with ordinary least squares, with standard errors clustered at the position level reported in parentheses.

Figure 7: Privacy Norms: Salary vs. Seniority

a. Is it Acceptable to Ask?

b. Are you Uncomfortable Asking?

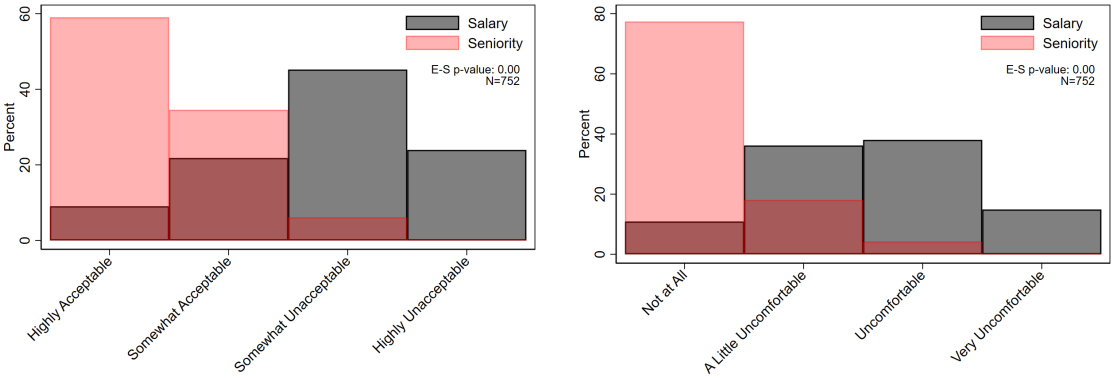

Notes: Panel (a) shows the distribution of responses to the question *Unacceptable*, asking whether it is "socially acceptable to ask someone about their salary/seniority". Panel (b) shows the distribution of responses to the question *Uncomfortable*, eliciting whether the respondent finds it "uncomfortable to ask information about salary/seniority to your peers."

Table 1: Randomization Balance Test

	All	Survey Type			Reward Size					
		Salary	Seniority	P-value	$13	$26	$39	$52	$65	P-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female (=1)	0.72	0.72	0.73	0.57	0.70	0.72	0.73	0.73	0.75	0.89
	(0.02)	(0.02)	(0.02)		(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	
Age (Years)	28.72	28.69	28.76	0.82	28.70	28.67	28.36	28.50	29.32	0.40
	(0.17)	(0.24)	(0.23)		(0.40)	(0.35)	(0.37)	(0.37)	(0.38)	
College (=1)	0.86	0.87	0.86	0.65	0.85	0.87	0.86	0.85	0.89	0.82
	(0.01)	(0.02)	(0.02)		(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	
Seniority (Years)	4.22	4.13	4.31	0.50	4.27	4.14	4.10	4.23	4.36	0.97
	(0.13)	(0.18)	(0.19)		(0.30)	(0.27)	(0.31)	(0.28)	(0.29)	
Own Salary (Masked)	1.00	1.03	0.97	0.24	0.96	1.01	0.95	1.04	1.03	0.58
	(0.02)	(0.03)	(0.03)		(0.05)	(0.05)	(0.04)	(0.05)	(0.05)	
Observations	752	376	376		135	160	131	165	161	

Notes: Average pre-treatment characteristics of the employees, with standard errors in parentheses. *Female* takes the value 1 if the employee is female and 0 otherwise. *Age* is the employee's age (in years) as of December 2017. *College* takes the value 1 if the employee finished College or a higher degree, and 0 otherwise. *Seniority* is the number of years from the date when the employee joined the company until December 2017. *Own Salary* is the employee base monthly salary as of December 2017 (due to the sensitive nature of the data, we do not reveal the unit of measurement for this variable). Column (1) corresponds to the entire subject pool. Columns (2) and (3) break down the sample in the two treatment groups that subjects were randomly assigned to: the survey about salary or about seniority, with column (4) showing the p-value of the null hypothesis that the averages are the same across these two groups. Columns (5) through (9) break down the sample in the five treatment groups regarding the reward amount, with column (10) showing the p-value of the null hypothesis that the averages are the same across these five groups.

Table 2: Comparison Between the Two Survey Types: Salary Vs. Seniority

	Abs. Error (Std.)		WTP ($)	
	Naive	Actual	Signal	Privacy
	(1)	(2)	(3)	(4)
Seniority Dummy (=1)	-0.020	-0.544***	-48.597**	-38.932***
	(0.052)	(0.096)	(21.973)	(7.636)
Constant	0.724***	1.069***	178.770***	67.146***
	(0.035)	(0.091)	(18.503)	(5.761)

Notes: N=752. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different regression and based on a different dependent variable: the standardized mean absolute error if the respondent had hypothetically entered her own salary/seniority as her guess (column (1)),the standardized mean absolute error of the actual guess provided by the respondent (column (2)), the willingness to pay for a signal of the average salary/seniority among five peers (column (3)), the willingness to accept or pay for sending an email revealing the respondent's salary/seniority to five of his or her peers (column (4)). The right hand size variable, *Seniority*, equals to 1 if the respondent was assigned to the survey about seniority and 0 if the respondent was assigned to the survey about salary. Columns (1) and (2) are estimated with Ordinary Least Squares, while columns (3) and (4) are estimated with an interval regression model. The standardized mean absolute error is defined as the difference between the individual's guess for the average salary/seniority in the sample of five peers and the actual average of those five peers, divided by the standard deviation of salary/seniority within the peer group.

Table 3: Salary Information: Heterogeneity by Gender

	Attitudes			Error		Absolute Error	
	Uncomf.	Unacc.	Recipr.	Actual	Naive	Actual	Naive
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female (=1)	-0.077	0.132	0.001	-0.006	0.013	0.003	-0.009
	(0.098)	(0.101)	(0.033)	(0.024)	(0.026)	(0.017)	(0.020)
Constant	2.486***	2.065***	0.888***	0.003	0.019	0.155***	0.142***
	(0.082)	(0.082)	(0.028)	(0.023)	(0.022)	(0.016)	(0.016)
	Accuracy (pp)	Perceived Acc. (pp)		Extra Week (pp)		WTP ($)	
	Actual	Direct	Indirect	Direct	Indirect	Signal	Privacy
	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Female (=1)	0.071	-0.062	-0.112**	-0.022	0.021	19.047	-8.711
	(0.048)	(0.043)	(0.052)	(0.028)	(0.033)	(35.835)	(8.458)
Constant	0.215***	0.603***	0.508***	0.599***	0.243***	165.143***	57.687***
	(0.041)	(0.039)	(0.043)	(0.039)	(0.037)	(30.905)	(7.476)

Notes: N=376. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. Each column corresponds to a different regression and based on a different dependent variable: the survey measures *Unacceptable* (column (1)), *Uncomfortable* (column (2)) and *Reciprocal* (column (3)), the error of the actual guess provided by the respondent (column (4)), the error if the respondent had hypothetically entered his or her own salary as guess (column (5)), the previous two outcomes but for absolute error instead of error (columns (6) and (7)), a dummy variable that takes the value 1 if the individual won the guessing game (column (8)), the expected probability to win the game without the extra week (columns (9) and (10), based on self-reported and incentive-compatible measures respectively), the expected gain in probability of wining the guessing game with the extra week (columns (11) and (12), based on self-reported and incentive-compatible measures respectively), the willingness to pay for a signal of the average salary among five peers (column (13)), and the willingness to accept or pay for sending an email revealing the respondent's salary to five of his or her peers (column (14)). The right hand size variable, *Female*, equals to 1 if the respondent is female and 0 if male. Columns (11) and (12) control for the probability of winning the game without the extra week. All columns are estimated with Ordinary Least Squares, except for columns (13) and (14) which are estimated by means of an interval regression model.

Online Appendix (For Online Publication Only)

A Further Details and Analysis

A.1 Characteristics of Subject Pool

We present descriptive statistics about the employee data in Table A.1. Column (5) corresponds to the final sample of 752 survey respondents. Column (1) corresponds to the universe of employees. The comparison between columns (1) and (5) implies that our sample is quite representative of the universe of employees. While some of the differences in gender, age, education and seniority are statistically significant, they are always economically small. For instance, the subject pool is 73% female vs. 71% female in the universe, the mean ages are 28.7 vs. 30.5 years old, the shares of college graduates are 85% for both, and the mean seniority is 4.2 vs. 4.7 years.

The only noticeable difference between columns (1) and (5) of Table A.1 is with respect to salary: our subject pool is 30% poorer than the universe of employees. The reason for this difference in average salary is quite simple: we did not send the survey invitation to employees in the highest paybands. These excluded employees, such as the CEO and vice-presidents, have salaries that drive up the average salary in the universe of employees quite a bit. To demonstrate this, columns (2) and (3) provide summary statistics for the sample of individuals who were not invited and were invited to the survey, respectively. The comparison of average salary across these two columns show that the bulk of the difference in mean salary between the subject pool and the universe of employees is coming from the selection of employees to be invited to the survey. For the sake of completeness, columns (4) and (5) provide statistics for employees who were invited to the survey but did not respond and individuals who responded, respectively. The average salary of the survey respondents is similar to the average salary of non-respondents.

Table A.1: Characteristics of Subject Pool

	All	Invited		Responded	
		No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)
Female (=1)	0.71	0.69	0.76	0.79	0.73
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Age (Years)	30.45	30.90	28.98	29.15	28.72
	(0.06)	(0.07)	(0.11)	(0.14)	(0.17)
College (=1)	0.85	0.85	0.85	0.85	0.85
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Seniority (Years)	4.75	4.85	4.44	4.59	4.22
	(0.04)	(0.05)	(0.08)	(0.11)	(0.13)
Own Salary (Masked)	1.42	1.55	1.00	1.00	1.00
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Observations	(Masked)	(Masked)	1,899	1,147	752

Notes: Average pre-treatment characteristics of the employees, with standard errors in parentheses. *Female* takes the value 1 if the employee is female and 0 otherwise. *Age* is the employee's age (in years) as of December 2017. *College* takes the value 1 if the employee finished College or a higher degree, and 0 otherwise. *Seniority* is the number of years from the date when the employee joined the company until December 2017. *Own Salary* is the employee base monthly salary as of December 2017 (due to the sensitive nature of the data, we do not reveal the unit of measurement for this variable). Column (1) corresponds to the entire company. Columns (2) and (3) break down the universe of employees by those who were not invited to participate in the survey and those who were invited, respectively. Columns (4) and (5) break down the employees who were invited to the survey by those who did not complete the survey and those who did, respectively. We do not reveal the total number of employees to protect the identity of the firm.

B Email with Invitation to Survey

Dear [Employee's Full Name],

We would like to invite you to participate in a survey. It takes less than 30 minutes to complete it. In the survey, you will participate in games which will give you the opportunity to earn rewards of $13 and upwards.

This survey is conducted by [Bank's Name] in collaboration with researchers from universities from the United States. Your responses will help us understand how to better communicate with our employees.

Follow this link to take the survey

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

Any rewards earned in the survey will be deposited automatically in your bank account within two business days of survey completion. Should you have any inquiries about your rewards, please contact us at [Survey Team's Extension Number].

In case of technical problems with the survey, please contact IT Support ([Extension Number]).

You were selected at random to receive this invitation. All your survey responses will remain confidential.

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

Sincerely,

Chief Economist,
[Bank's Name]

C Survey Instrument

Dear [Respondent's Name],

You are invited to participate in a survey conducted by [Bank Name]. This survey was designed in collaboration with academic researchers from universities in the United States. This survey will teach us about how [Bank name]'s employees learn about their workplace, earnings, and career prospects.

In the survey, you will participate in games which will give you the opportunity to earn rewards of $13 and upwards. Any rewards earned in the games will be deposited in your bank account automatically.

ALL SURVEY RESPONSES ARE COMPLETELY CONFIDENTIAL.

Thank you in advance for your participation!

Sincerely,

Chief Economist
[Bank Name]

○ Please click here to confirm that you are [Respondent's Name] and you would like to take part in this study

>>

You'll play a series of short games. In these games, we will ask how much you would pay for something, or how much you would sell something for. Then a computer will bid against you.

Next to these questions, you will see the message: "You are bidding against a computer, not a person, **it is best for you to report truthfully**." Here's the explanation:

These games are just 'pretend' but we will choose **a few lucky participants** from this survey to play for real!

So let's say we ask 'How much would you be willing to pay for an iPhone X? If you say $1,000, and the computer says $800, we will give you the iPhone X for free. If you say $800 and the computer says $1,000, we will give you $1,000. This auction was designed by economists so that it is best for you to say your true preference: that is, say exactly how much you would really be willing **to** pay for the Iphone X.

In all of these games, you can earn money, but you will never lose money.

Remember that this is not a regular first-price auction, in which it is optimal to bid less than your true valuation. In this type of auction, called the **second-price auction**, you will be always worse off if you try to under bid or over bid your true valuation.

>>

Let’s do a practice question: How much would you be willing to pay for an iPhone X?

Remember you are bidding against the computer, so it is in your own interest to report truthfully. Also, if you don't want an iPhone X, you can always sell it (for your reference, the market price of the iPhone X is around $1,000).

0 $

>>

0259

Let's play a guessing game. **You have 3 minutes to answer this question, or you won't qualify for the prize.**

The game consists of guessing the average [Salary/Seniority] among five of your peers. Your peers are defined as your coworkers who work in your same position (Teller) and unit (Branch 25). According to our records, you have 12 peers in this group. The question is about the following sample of 5 out of the 12 peers:

[First and Last Name of Peer 1]
[First and Last Name of Peer 2]
[First and Last Name of Peer 3]
[First and Last Name of Peer 4]
[First and Last Name of Peer 5]

The rules of the guessing game are simple: if your guess falls within +/-5% of the true average among these 5 peers, you will receive a reward of $[Reward Amount].

What is the average [Salary/Seniority] among the 5 selected peers as of December 2017?

0 [$/Years]

>>

You have the opportunity to buy the following information: the average [Salary/Seniority] among a *different* random sample of 5 of your 12 peers.

If you buy this information, you will be given the opportunity to use that information to revise your guess. This can improve your chances of winning the $[Reward Amount] reward.

What is the maximum amount of money you would be willing to pay for the information (the average [Salary/Seniority] among a random sample of 5 peers)?

Remember that you are bidding against the computer, so it is in your own interest to report truthfully.

0 $

>>

You have not been selected to be able to buy information. Please continue with the survey.

>>

Before proceeding with the survey, we want to introduce a new type of question.

Let's say you are playing the following game: a coin is flipped, and you have to guess whether it fell head up or tails up. If you guess correctly, then you get $100; if you do not guess correctly, you get $0. Note that the expected prize for this game is $50.

How much should we pay so that you give up the right to play this coin game? Remember that you are bidding against the computer, so it is in your own interest to report truthfully.

0 $

>>

You have entered a guess of $[Guess] about the average [Salary/Seniority] of your peers. If this guess falls within +/-5% of the truth, you will be rewarded $[Reward Amount]. What do you think is the probability that your guess will fall within +/-5% of the truth?

0 %

>>

Since you expect to win a reward of $[Reward Amount] with a probability of [Probability]%, your expected reward is $[Probability*Reward Amount].

Now, we want to offer you a fixed amount of money to not play this guessing game. We do this with all participants, regardless of their guesses. What is the smallest amount of money that you would be willing to accept to give up the right to play the guessing game? Remember that you are bidding against the computer, so it is in your own interest to report truthfully.

0 $

>>

It has been determined that you will continue to play the guessing game.

>>

Some participants, selected at random, will be given an additional week to revise their guesses. You will find out if you get the additional week in the next screens.

Please imagine that you were given the additional week. You could use this additional week to search for information and increase the accuracy of your guess. For instance, you could use this additional time to ask some of your peers about their [Salary/Seniority].

Remember that, without the additional week, you expect a probability of winning the guessing game of [Probability]%. If you had the additional week, what is the probability that you would win the guessing game?

0 %

>>

Now, imagine that you were given the additional week to revise your guess, which you can use to search for information and increase the accuracy of your guess. Since you expect to win a reward of $[Reward Amount] with probability [Probability]%, that means that your expected reward with the additional week is $[Probability*Reward Amount].

Without the additional week, you were willing to accept $[Amount Entered] in exchange to give up the right to play the guessing game. Now, with the additional week, how much should we pay so that you give up the right to play this game? Remember that you are bidding against the computer, so it is in your own interest to report truthfully.

0 $

>>

You were not chosen to have the extra week to revise your guess.

>>

The survey is almost over. In this last section, we want to know how you would feel if someone revealed information about your [Salary/Seniority] to some of your peers.

We could send an email to your 5 selected peers revealing your [Salary/Seniority], including your full name. This email would explain that this message was sent in the context of a game.

We will leave it entirely up to you whether we send this email or not.

Would you want us to send this email to your peers?

- ○ Yes
- ○ No

>>

Since you do not want us to send this email, we want to offer you some money in exchange of sending this email out.

Please let us know the minimum amount of money that you would accept to let us send this email to your peers? Remember it is in your best interest to answer truthfully. If you enter an amount below $100, there is a chance that we send the email and you get the amount you entered as compensation. However, any amount above $100 would surely result in no email being sent (but also no compensation).

Remember that your response will be compared to the bid from the computer, so it is in your own interest to report truthfully.

0 $

>>

Do you find it uncomfortable to ask information about [Salary/Seniority] to your peers?

- Not at all
- A little uncomfortable
- Uncomfortable
- Very uncomfortable

If you ask a peer about his or her [Salary/Seniority], would you expect this peer to ask you about your [Salary/Seniority]?

- Yes
- No

Is it socially acceptable to ask someone about their [Salary/Seniority]?

- Highly unacceptable
- Somewhat unacceptable
- Somewhat acceptable
- Highly acceptable

>>