THE SALARY TABOO:
PRIVACY NORMS AND THE DIFFUSION OF INFORMATION

Zoë B. Cullen
Ricardo Perez-Truglia

The Salary Taboo: Privacy Norms and the Diffusion of Information
Zoë B. Cullen and Ricardo Perez-Truglia
NBER Working Paper No. 25145
October 2018, Revised November 2019
JEL No. C93,D83,D91,J3,J71,M5,Z1

ABSTRACT

The diffusion of salary information has important implications for labor markets, such as 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 755 employees at a multibillion-dollar corporation to study how people search for and share salary information. We show that most employees have inaccurate beliefs about the average salary of their peers. We provide evidence that such misperceptions arise, in part, due to search costs, and we provide suggestive evidence that these costs are associated with privacy norms. Last, we show that, contrary to widespread belief, there are no significant gender differences in misperceptions, search costs, and privacy norms.

Zoë B. Cullen
60 N. Harvard, Rock Center 210
Harvard Business School
Boston, MA 02163
zcullen@hbs.edu

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

1. INTRODUCTION

Employees rely on information about salaries in several ways. For example, they may need the information to negotiate raises or to switch positions or employers. Most employers provide limited information about salaries. Thus, employees' knowledge about salaries depends largely on their ability to communicate with each other. There is a widespread belief that access to information about salaries is imperfect and unequal. However, there is no direct evidence on the magnitude and sources of these frictions. 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 being informed about salaries, but that information has a cost. Employees must spend time and energy searching for salary information, and they may face other social costs for inquiring about sensitive data. Some signs suggest that these information frictions are significant. For example, while most employees desire to be better informed about the salaries of their coworkers, they rarely discuss salaries with their coworkers (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). Indeed, alleged information frictions have inspired several policies over the past decades, 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 magnitude and sources of these information frictions.

We collaborated with a large firm to conduct a field experiment with its employees. We designed an incentivized survey that, combined with rich administrative data from the same organization, allows us to address some of the key questions on how employees learn about salaries at the workplace. Our research design allows us to assess employees' perceptions of peer salaries, employee confidence in their beliefs about peer salaries, and the roles of search

---

[1]There is no consensus on 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).

costs and privacy norms in generating these misperceptions.

First, our survey measures perceptions of peer salaries by asking employees to guess the average salary of a random sample of five of their peers (i.e., coworkers in the same unit and with the same position title). For example, a Junior Researcher from Investment Banking is asked to guess the average salary of five other Junior Researchers working in Investment Banking, listed by first and last name. To provide incentives for truth telling, employees whose guesses fall within 5% of the true average salary receive a monetary reward ranging from \$13 to \$63.[2] By comparing the employees' guesses to the the actual salaries from the administrative records of the firm we can measure misperceptions about peer salaries.

Second, our survey measures how confident employees are in their beliefs about the salaries of others. After announcing the prize associated with accurately guessing the salaries of five peers, we elicit how much money the participant would accept to forgo the game using our incentive-compatible methods. We also elicit this belief in a non-incentivized manner by asking participants directly how likely it is that they provided accurate guesses.

Third, our survey measures the role of search costs in explaining salary misperceptions. We offer respondents the opportunity to have an extra week to search for salary information on their own, which allows them to improve their guesses and thus increase the chances of winning the guessing game. For example, participants could use the extra time to ask one or more of the five peers in the list about how much they make or to research information by other means. We elicit the probabilities of winning the guessing game with and without the additional week using an incentive-compatible method. We use the difference between these probabilities as a measure of the employees' willingness to search for information on their own.

To validate this measure of willingness to search for information, we randomize the size of the rewards for guessing correctly to one of five different values from \$13 to \$63. This variation allows us to test the hypothesis that individuals are more willing to search for information when they stand to gain more from it (Woodford, 2001; Sims, 2003; Mankiw and Reis, 2002; Reis, 2006).

As a complementary test of search costs, following Cullen and Perez-Truglia (2018), we measure the willingness to pay for salary information. We provide subjects with the opportunity to acquire a readily available signal about the true average peer salary from the experimenter. Using a standard incentive-compatible method, we elicit the willingness to pay for this signal. If individuals find it costly to search for information on their own "in the wild," they should be willing to pay significant amounts for this signal.

---

[2]These amounts, as well as all other monetary amounts discussed in the paper, have been transformed In United States dollars using PPP-adjusted exchange rates from February 2018.

Finally, we use multiple strategies to explore the role of privacy norms in information frictions. We include subjective questions to measure social norms around privacy, such as whether it is socially acceptable to ask peers about their salaries or whether subjects feel comfortable asking others about their salaries. To address the usual concerns with subjective data, we also designed a revealed-preference measure of these preferences for privacy. We offer the subject the opportunity to reveal his or her own salary with five of his or her peers, in a verifiable way. Using an incentive-compatible method, we elicit the subjects' willingness to reveal (or conceal) this information.

To further assess the role of privacy norms, we study how individuals learn in- formation about the seniority of their peers (instead of learning about salary information). Like salary information, seniority information can be useful in making important career choices. However, unlike salary information, seniority information may not be subject to strong privacy norms. We randomized subjects into one of two survey types: the salary survey or the seniority survey. The seniority survey is identical to the salary survey except that instead of revolving around the average *salary* of peers, it revolves around their average *seniority*. This randomization allows us to test for differences in privacy norms between salary and seniority information and whether they translate into differences in willingness to search for information.

We conducted the experiment with a sample of 755 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 provides a context that seems most relevant to study information frictions. For example, the firm does not disclose salary information, employees reportedly desire more salary transparency, and employees rarely discuss salaries with their coworkers. Several studies show that these three facts are common in organizations from several countries, including but not limited to the United States (Trachtman, 1999; Edwards, 2005; Hegewisch et al., 2011; Glassdoor, 2016; PayScale, 2018).

The results from the experiment suggest that employees have significant misperceptions about the salaries of their peers, that those misperceptions are partly due to information search costs, and that those costs are in part due to privacy norms.

We start by showing a couple of results that, despite some methodological differences, replicate the results from an earlier study (Cullen and Perez-Truglia, 2018) that motivated this work. We show that the guesses about the average salary of the five peers have a mean absolute error of 15%. This level of misperception is what we would expect if employees would have used their own salary as a guess for the salary of their peers, thus suggesting that employees have little information about salaries besides their own salary. Also consistent with our earlier study, we find that many employees demonstrate significant willingness to

pay for information about the salaries of their peers. That evidence suggests that some employees are not misinformed due to lack of interest. For example, among the top half of the sample, the willingness to pay for salary information has a median of $130 and a mean of $369 (approximately 1 and 3 weeks' worth of salary, respectively)[3]

Next, we provide a novel set of results that has not been previously documented. We show that employees are partially aware that they do not have perfectly accurate beliefs. However, on average, they are overconfident about their accuracy. Based on our incentivized questions about the perceived likelihood of winning, the average employee believes she or he has a 43% probability of guessing within five percentage points of the truth, while in reality the fraction of accurate responses was 32%. This discrepancy is even larger when we use the non-incentivized measure of accuracy, as employees' direct reports about the likelihood of winning is on average 56%. This evidence suggests that overconfidence about employees' current knowledge of salaries may make them under-invest in gathering salary information.

We provide evidence that search costs also play a role in explaining the salary misperceptions. When presented with financial incentives to do so, some individuals are willing to search for information on their own. Indeed, consistent with the rational inattention hypotheses, employees who are randomly assigned to higher game rewards should search more intensively. However, many employees are not willing to search even when provided significant incentives. For example, 25% of respondents are not willing to search for information on their own, even when offered the highest prize for accuracy ($63).

Next, we discuss the role of privacy norms as a source of search costs. According to self-reported data, there are privacy norms against discussing salaries with coworkers: 69% of employees think that it is socially unacceptable to ask coworkers about their salaries, and 89% would feel uncomfortable asking this information. Moreover, the revealed-preference measure also suggests that most employees prefer to keep their salary information private. Whereas a minority of employees (20%) prefer to share personal salary information with five of their peers, most (80%) prefer to conceal this information. Moreover, 38% are not willing to reveal their salary information even if offered $125 to do so. We show that those who feel more uncomfortable asking others about their salaries are less willing to search for salary information on their own. This evidence suggests that privacy norms have real effects on the decision to acquire information.

To further explore the link between privacy and search costs, we turn to the comparison between seniority and salary information. We find that, as expected, seniority is a less sensitive topic than salary: whereas 69% of employees find it unacceptable to ask coworkers about

---

[3]The bottom half of the sample is willing to pay less than $13 for the signal about the salary of peers, suggesting that they may be misinformed about the salaries of their peers primarily due to a lack of interest.

their salary, only 6% find it unacceptable to ask coworkers about their seniority. Moreover, the subjective data align with the revealed-preference data: on average, employees are willing to reveal their salary to five peers for $51 but willing to share their seniority for just $23.

We show three pieces of evidence that the more sensitive information (salary) does not flow in the network as much as the less sensitive information (seniority). We show that employees are as accurate guessing salaries as they would be if they just reported their own salaries. However, employees are substantially more accurate when guessing seniority than they would be if they just reported their own seniority. This evidence suggests that employees have access to other information about seniority besides their own seniority but have no other information about salaries other than their own.

Also, we show that employees who are more central in the network or who are more connected to specific peers are better informed about the seniority of those peers. In the presence of social learning, we would expect individuals who are most central in a network to be the ones who are best informed, as past theory and evidence has shown (Alatas et al., 2016; Banerjee et al., 2013), and indeed this is what we find with respect to seniority information. In contrast to this, the accuracy of beliefs about salaries is orthogonal to their centrality or their connectedness to peers, suggesting little social learning when it comes to salary information.

Although social norms may explain part of this demand for privacy, secrecy also may be strategic in nature. For example, some individuals may not want to reveal to others that they are being under-paid, because it would damage their social image. Others may not want to reveal that they are being over-paid, because that may generate resentment from their coworkers or undermine their leverage in future salary negotiations. To assess this strategic role, we look at the relationship between the demand for privacy and the employee's position in the distribution of salaries. We find that the second channel dominates: those who perceive themselves to be higher relative earners are less willing to reveal their salaries.

As aforementioned, there is widespread belief that access to salary information is not only imperfect, but also *unequal*. In particular, there is a widespread belief that pay secrecy and privacy norms disproportionately hurt women (Babcock and Laschever, 2009). For example, survey data indicates that women feel less informed than men about the salaries where they work (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 have slightly more accurate perceptions than their male counterparts. Moreover, we find that the gender differences across other outcomes also are small, statistically insignifi-

cant, and precisely estimated: female and male employees are equally willing to say that they are uncomfortable asking others for their salaries, equally willing to search for information, equally willing to pay for readily available information, and equally willing to reveal their own salaries to their 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 et al., 2017; Baker et al., 2019). Conlon et al. (2018) find that there is a demand for information about the wages employees could command from outside job options. Their paper reports that a non-college-educated employee would pay $175 per year to have complete information about outside options, while a college-educated employee would give up $817 every year for the same information. Yet, there is little *direct* evidence on the magnitude and sources of frictions preventing people from gathering salary information.[4]

This study builds on our previous work documenting significant misperceptions of peer and manager salaries (Cullen and Perez-Truglia, 2018). Our earlier work provides evidence that employees have misperceptions about the salaries of their peers and shows that, even when provided with information through an experiment, the information given to employees does not seem to travel through the peer network. We expand on this earlier finding by measuring the sources of the information frictions, with special emphasis on privacy norms.

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, Phan, and Szeidl (2015) recruited college students to play a "treasure hunt" game in which they earned rewards by collecting information from peers. Other studies exploit natural incentives for information diffusion. For example, Beaman, Dillon, and Lori Beaman (2018) seeded useful information about composting and measured its diffusion in an agricultural network. These papers provide evidence that, even in settings where information is mutually beneficial, its diffusion is highly imperfect. Chandrasekhar, Golub, and Yang (2018) investigate how individuals may avoid seeking helpful information when doing so would signal low skill or induce shame. Our contribution to this literature is twofold. First, we contribute a new method to measure the willingness

---

[4]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 et al., 2012; Perez-Truglia, 2019; Mas, 2016, 2017; Breza et al., 2018; Cullen and Pakzad-Hurson, 2017; Cullen and Perez-Truglia, 2018).

to search for information and the willingness to share information with others. Second, we show evidence that the diffusion of information may be worse when the information to be disseminated is subject to privacy norms.

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, Catalini, and Tucker (2017) and Adjerid, Acquisti, Brandimarte, and Loewenstein (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 information frictions hurt minorities disproportionately (Phillips, 2009; Colella et al., 2007). 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, which could in itself have an effect on salary negotiations.

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

2. RESEARCH DESIGN

2.1 *Conceptual Framework*

The research design is inspired by a simple model of endogenous information acquisition. Assume that employees care about their coworkers' average salary; this could be useful for all sorts of things, such as negotiating salary or to deciding whether to change jobs or employers. More specifically, we assume the employee's utility is a function of the accuracy of his or her belief about peer salaries.

Employees can take actions in order to improve the accuracy of their peer salary beliefs. While those actions would increase their utility, they come at a cost – e.g., the employees must spend time and attention to acquire the information, plus they may need to incur the cost of inquiring about sensitive information. As a result, the decision of how much

information to acquire boils down to a cost-benefit analysis. Figure 1.a provides a graphical representation of this cost-benefit analysis. The x-axis represents the employee's accuracy. 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.

Figure 1.a shows the Marginal Benefit (MB) and Marginal Cost (MC) curves, under the classical assumption that the MC curve has an upward slope and the MB curve has a downward slope.[5] The point q_0 on the x-axis represents the point at which a rational individual stops acquiring new information: i.e., exactly at the point where the marginal cost of information equals the marginal benefit of information. We will come back to this framework to motivate and illustrate the different survey measurements that follow.

2.2 *Survey Types: Salary Versus Seniority*

Appendix C includes a full sample of the survey instrument. One of the key aspects of the survey is that participants are assigned with equal probability to one of two survey types:

- *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 from a specific branch would 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.

---

[5]The assumptions about the slopes of the MC and MB curves help simplify the explanation, but they are not crucial for the following results. Moreover, there are simple ways of micro-founding these slopes. For example, a given employee may be more comfortable asking some peers than others. This employee will by asking the peer he or she feel most comfortable asking, then move to the second most comfortable, and so on. This will result in an upward slope for the MC curve. In turn, the downward slope of the MB curve can arise due to the principle of diminishing marginal returns: a given increase in accuracy may be more useful to an individual who is completely uncertain about the topic than to an individual who is almost certain.

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" units. To simplify exposition, we only describe the version asking about salary in detail.

Just like information on peer salary, the information on 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, there may be stronger privacy norms around salary information than around seniority. For example, while there are countless studies mentioning the term "salary taboo," there are no mentions of a "seniority taboo." As a result, the comparison between the results of the salary and seniority surveys may give us hints about the role of privacy norms for the diffusion of information.

2.3 *Incentive-Compatible Elicitations*

In this study we strive for incentive-compatible survey methods whenever possible. While incentivized surveys are generally welcomed in Economics, this seems to be particularly valuable in the topic of privacy. For example, individuals tend to *say* that they value their privacy a lot, but then their behavior reveals that they do not value it nearly as much as they say. Athey et al. (2017) documents a 54% decline in the likelihood that an MIT undergraduate protects their friends' contact information when they introduce a small incentive, free pizza. This is despite the fact that respondents rank friends' contact information as the second most private piece of data, just below social security numbers, in the National Cyber Security Alliance (NCSA) survey. When asked directly, 60% stated they would never feel comfortable sharing these contact details if asked.

To elicit valuations (for information, privacy and other things) in an incentive-compatible way, we employ the traditional Becker-DeGroot-Marschak (BDM) method. We use the 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 makes it a dominant strategy for respondents to 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 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 order to ensure BDM is incentive compatible, the subjects must know that the probability is positive, but it does not matter exactly what the value of the probability is. Consistent with this view, it is well documented that it normally does not matter the exact probability with which these mechanisms are executed as long as the probabilities are greater than zero (Cason and Plott, 2014). 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 distracting 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, 2019; Fuster et al., 2018), which differs from another common mechanism in which subjects must pay out of their pockets. We did not implement this latter mechanism because the firm wanted to avoid collecting payments from its employees' bank accounts.

While the BDM method has some advantages, it is of course not perfect. Some of their imperfections have been documented in the literature. For example, some subjects may shade their valuations, as if they were playing a first-price auction, even though that is a dominated strategy (Cason and Plott, 2014). In some special cases, it may not even be incentive-compatible to report one's true valuation.[6]

One of the ways in which we try to mitigate these sources of biases and measurement error is by including a training module at the beginning of the survey (Cason and Plott, 2014). 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 the BDM elicitation.[7] Additionally, in order to show that there is some signal among

---

[6]This situation can arise if the auctioned object is a lottery (Karni and Safra, 1987), and even if the auctioned object is non-random (Horowitz, 2006)

[7]One of the training questions elicits the willingness to pay for an iPhone X. The other training question elicits the willingness to forfeit a lottery that pays $100 with probability 50%. There are no correct or

the noise, we present a number of validation exercises.

Despite these efforts, measurement error may exist. As a result, while subjects may seem quite heterogeneous in their bids, some of that heterogeneity may simply reflect measurement error. Fortunately, when we are looking at the association between variables, the measurement error will generate attenuation bias: i.e., making those associations look less strong than they actually are.

2.4 *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 peer groups can have between 6 to 89 members, with a median size of 21 employees. To make the guessing game more comparable across participants with different peer group sizes, 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 chosen peers.

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 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. This was intended to make sure that participants did not have time to search for information if they wanted to, because we want them to give them the opportunity to search later in the survey.

2.5 *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. Additionally, we implement an alternative measure that is incentive-compatible.

---

incorrect answers for these questions. However, if the subjects provide responses that are abnormally high or abnormally low, we provide some feedback explaining describing the typical responses to these questions.

This measure is illustrated in Figure 1.b. This is the situation of an individual who has searched up to the point q_0, where the marginal benefit of searching equals the marginal cost. Now imagine that the individual is given the unexpected opportunity to participate in the guessing game described above, in which 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 the curve upward by exactly the amount of the accuracy reward. Since the individual is not given any time to search for information before providing, they must remain at q_0. In this case, playing the guessing game increases the welfare of the individual 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: i.e., the certainty equivalent that would make the subject forfeit the right to play the guessing game.

In the survey, we elicit that certainty equivalent using the BDM method. Subjects can ask for any dollar value from 0 up to the full reward amount. We explicitly mention that, due to the BDM mechanism, it is in the respondent's best interest to bid their true willingness to accept for giving up 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 can normalize the bid that the subject provides as a share of the reward 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 was risk-neutral, like we assumed in Figure 1.b, this ratio would reveal the perceived probability of winning the guessing game.

Despite the advantage of being incentive compatible, this alternative measure has some disadvantages. First, it may introduce measurement error because the willingness to forfeit the game depends not only on the perceived probability of winning but also on risk preferences. Second, it may introduce measurement error because the question is harder to understand than the simple probability question.[8]

2.6 *Willingness to Search for Information*

We want to elicit whether the respondents are willing to search for more information on their own. 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 even

---

[8]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.

consult with Human Resources.

We can use Figure 1.b to illustrate this measure. This figure corresponds to the introduction of the guessing game, which shifted the MB curve from MB to MB'. Since the individual was not given enough time to search for information before providing his or her guess, the individual must stay at q_0. However, if the individual had extra time to search for information before providing the guess, he or she would want to search for additional information up to point q_1, where 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.

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 subjects 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).

One potential 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. We illustrate this alternative measure in Figure 1.b.. The shaded triangle between q_0 and q_1 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 tell the participants that some subjects may get the opportunity to buy the additional week, and ask them to bid for this opportunity with the usual BDM mechanism. Again, we remind them that it is in their best interest to bid their true valuation. And also 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 we can normalize this measure a share of the total game reward, which in the case of risk-neutral individuals should be equal $q_1 - q_0$.

The randomization of the size of the accuracy rewards give us a rational inattention hypothesis. This is illustrated in Figure 1.c, which 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 other words, we can test if individuals who were randomly allocated to higher reward amounts were indeed more willing to search for information.

14

2.7 *Willingness to Pay for Readily Available Information*

One limitation with the measure of willingness to search is that it gives a sense of what are the net befits from acquiring more information, but it is uninformative about the gross value of information. Figure 1.d illustrates this point by depicting two 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 information values. 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, the 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]

In the survey, after subjects enter their guesses, they are given the chance to buy a piece of information, that is, a signal related to the peer salary they are trying to guess. This signal consists of the average salary or seniority among a different sample of five peers (i.e., from a draw of five peers). Although not perfectly informative, 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.

Note that the reward randomization again gives us a test of rational inattention. 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

---

[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.

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.

2.8 *Demand for Privacy and Privacy Norms*

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.

The survey tells respondents that the experimenters are considering sending an email to five peers, implicitly the same five peers whose salaries or seniority they guessed. This email will include 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.[11] 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 take 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.

At the end of the survey, we included 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."

The last question is intended to measure a reciprocity norm. Individuals may be averse to asking about salary and seniority not because they want to avoid bothering others, but

---

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

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."

3. 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., 2019).

The firm is typical in other relevant respects. It does not have an open salary policy. 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 and less than a quarter disclose data on salary ranges (Scott, 2003). And a survey of employees from eight developed countries show that they are uninformed about salaries and want employers to be more transparent regarding pay (Glassdoor, 2016). Moreover, the standard employment contract at this firm 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 2001 survey of U.S. employees finds that more than one-third are forbidden from discussing their pay with coworkers (Day, 2007; Vault, 2001).

Also, in our firm, survey data indicate that almost half of the employees never discuss their salaries with coworkers. 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 (Gee, 2017).[12] Last, as shown below in the results section, in this firm there seems to be a taboo around salary discussions. This taboo is believed to be present in a broad range of countries including but not limited to the United States (Edwards, 2005), Canada (Bierman and Gely, 2004) and Israel (Fox and Leshem, 2005).[13]

One of the sociological explanations for the salary taboo can explain why this taboo

---

[12]In a survey of 1,022 employees from the United Kingdom found that less than half (48%) discuss salaries with their peers (Burchell and Yagil, 1997).

[13]For example, Fox and Leshem (2005) present survey data indicating that most individuals in Israel report feeling highly uncomfortable when asked about their pay and other financial matters.

is so universal. According to this theory, pay can act as a signal of self-worth, and thus individuals do not feel comfortable discussing pay because it is equivalent to discussing self-worth (Trachtman, 1999). Thus, as long as individuals perceive that pay is somewhat related to the marginal product of labor, which is arguably the case in all market economies, there will be room for the salary taboo.

3.1 *Survey Implementation*

We start with the universe of thousands of employees. We focus on two specific units of the firm, with some added filters.[14] 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, 755 individuals finished the survey, corresponding to a 39.7% response rate.[15] The median respondent took 15 minutes to complete the survey.

3.2 *Descriptive Statistics and Randomization Balance*

The subject pool includes employees from 46 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 I presents some descriptive statistics about the subject pool. Column (1) corresponds to the entire sample of 755 survey respondents: 73% 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.

---

[14]For instance, we exclude employees from the highest step of the corporate ladder. And to avoid any contamination, we exclude employees who participated in a previous survey that was related to peers' salaries (Cullen and Perez-Truglia, 2018).

[15]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).

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 I, 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 I provide the corresponding balance test for this treatment arm. Again, the results are consistent with successful random assignment.

4. RESULTS

4.1 *Misperceptions*

Figure 2 shows misperceptions about average peer salary, comparing two different benchmarks in panels (a) and (b). This figure indicates that only 32% 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 14.6%. These misperceptions 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 -1.5% (p-value=0.184).

Only 32% of respondents provide a guess of the average peer salary that is close (i.e., within 5%) of their own salaries. This evidence suggests that employees use other information sources besides their own salaries to come up with their guesses. 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=0.13). If individuals report their own salaries as their guesses for average peer salary, the mean absolute error (16%) is only slightly higher than the mean absolute error of the actual guesses provided by the subjects (14.6%).

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 (10%, instead of 15% in reality). This finding shows that most misperceptions are not caused by asking about a specific subsample of peers.

As a robustness check, we can take advantage of the fact that we replicated some of the measurements made in a previous study (Cullen and Perez-Truglia, 2018), with a similar sample but with some methodological differences: we used a quadratic scoring rule to incentivize responses, we used smaller reward amounts, and we asked about the average salary among all peers instead of just a sample of five peers. Despite these differences, the results from the two experiments are quite consistent in magnitude: the MAE of peer average guesses are 14.6% in this survey, which is in the same order of magnitude as the 11.5% reported in Cullen and Perez-Truglia (2018).

4.2 *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 32% 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).[16] 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.[17]

---

[16]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.

[17]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 32% of guesses that actually fall within that range.

4.3 *Willingness to Search for Information in the Wild*

Figures 3.a and 3.b, in addition to the probability of winning the game without the extra week, also 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 reported probability of winning the game with and without the extra week, divided by the probability of losing the game without the extra week. This sample excludes 30 individuals who reported 100% confidence in their initial guesses and for whom there could be no gain in certainty by construction. 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 ex-ante 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.[18] In any case, our measure of anticipated gains is the relevant measure for the decision to search for information or not.

It must also be noted that our measure of willingness to search may overestimate the "natural" willingness to search for salary information. There are two reasons for this. First, 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

---

[18]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.

ordinary alternatives. Second, our survey may have tacitly lifted the non-disclosure policy for the duration of our study.[19]

To validate this measure of willingness to search, we can turn to the rational inattention hypothesis discussed in Section 2.7. According to this hypothesis, individuals should be more willing to search for information when the gains from doing so are higher: i.e., a higher reward in the guessing game should cause individuals to search more intensively. Figure 4.b illustrates this test, by comparing the willigness to search with the (randomly assigned) reward amounts. Consistent with rational inattention, individuals who are assigned to higher rewards expect to search more intensively for information. More precisely, a $50 increase in the guessing reward results in an expected accuracy increase of 7 percentage points. This difference is not only statistically significant, but also economically large: this 7 percentage points increase implies a 12.5% improvement relative to the average perceived accuracy rate of 56 percentage points.

4.4 *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 2.1, 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. This willingness to pay for information is an order of magnitude higher than the guessing game rewards. As a result, these individuals cannot be bidding for the information with the main goal of winning the guessing game. Instead, 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. As noted by (Stigler, 1962), it is not difficult to rationalize high valuations

---

[19]However, half of the employees report that they do discuss salaries with their coworkers. This evidence suggests that employees are not significantly concerned about the firm's secrecy policy.

for salary information. For instance, 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 up to 5% of the annual salary (over two weeks' worth of pay) for this information.

As discussed in Section 2.3, the BDM method is imperfect and thus subject to biases and measurement error. More specifically, one special concern is that our estimates of willingness to pay may be sensitive to the elicitation method. As a robustness check, we can take advantage of the fact that we replicated some of the measurements made in a previous study (Cullen and Perez-Truglia, 2018) with a different methodology. In that earlier study, we measured willingness to pay for information in a comparable sample, but used the price-list method instead of the open-ended method used in this experiment. As shown in the Appendix (Figure A.1), the distribution of willingness to pay is similar across the two elicitation methods. This finding suggests that, consistent with other studies (Brebner and Sonnemans, 2018), our measures of willingness to pay are consistent across different elicitation methods.

To validate this measure of willingness to pay for information, we can turn to the rational inattention hypothesis discussed in Section 2.7. It is plausible that subjects bidding close to the median ($13) are bidding primarily to improve the chances of winning the guessing game. According to the rational inattention hypothesis, the individuals who face higher rewards in the guessing game should be willing to pay more for the information, because they stand to gain more from it. Figure 5.b 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=0.030).

The last validation test consists in looking at the association between the WTP for information and the career stage of the employees, using administrative data. We look at three career milestones that may be related to the WTP for salary information. We expect that subjects value salary information the most when they are up for a promotion or salary renegotiation, for which they are eligible at most once per year. We also expect active family planning, in particular maternity leave, to reduce the likelihood of salary negotiations and therefore the value of salary information. The results are presented in Table II. Aligned with these predictions, We find a significant relationship between the WTP for salary information and the subsequent career outcomes. Columns (1) of Table II shows that those in the upper quartile of the WTP (>$100) are also 8.5 percentage points more likely to receive a promotion in the next three months, 61% higher than the rate of promotion among those unwilling to pay $100 for a signal about the salary of their peers. Similarly, Col. (2) shows the willingness to pay for salary information is predictive of receiving a raise in the three months after the experiment. We interpret this as evidence that individuals value salary information highly at certain times in their careers, and much less at other times. Col. (3) shows that

those individuals anticipating family leave in the next 3 months value salary information less than those who do not. The overall rate of family leave is low among participants, approximately one percent of people, but none from this group were willing to pay $100 for salary information.

4.5 *Willingness to Share Information with Peers*

The above 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 the own salary information with five peers. Roughly 20% of employees prefer the experimenter to reveal their salaries to peers, and the remaining 80% prefer to avoid sending the email. Figure 6.a shows that, both among individuals who want to reveal and conceal their salaries, there is quite a bit of variation in the strength of their preferences. Roughly 40% of subjects have weak preferences for privacy because that they are willing to pay less than $5 to reveal or conceal their salary information. The remaining 60%, however, show strong preferences: a whopping 40% is not willing to reveal their salaries for anything less than $125. A surprising result is that 4% of respondents are willing to pay over $125 to reveal their salaries. This may seem puzzling 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 without appearing that they are showing off.

The unwillingness to share one's own salary with coworkers may reflect a direct preference for privacy, but it also may respond to strategic incentives. There are two mechanisms that fit the strategic motive. On the one hand, 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. Thus, the higher the relative salary of the employee within the peer group, the stronger the preference should be to keep the salary information private. On the other hand, the models of social status (Frank, 1984; Bursztyn et al., 2018) make the opposite prediction: employees with higher relative salary should be more excited about revealing their salary, because that will be a boost to their social status.

Our unique data on willingness to pay for privacy allows a direct comparison of these two mechanisms with opposing predictions. Figure 7 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 \$30, which is equivalent to 50% of the standard deviation of this outcome (\$76). In contrast, the relationship is downward sloping and statistically insignificant, for the willigness to share seniority information.[20] 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 from the manager or even put their relative salary at risk.

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.

Our survey data also suggests one way in which the unwillingness to reveal one's salary may exacerbate the unwillingness to search for information. When asked in our survey, 89% of employees report that asking another employee about his or her salary will result being asked back for their own salaries. This reciprocity norm suggests that some individuals want to avoid asking others about their salaries because they do not want to be asked back about their own 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. For this reason, we generate even stronger evidence against status concerns than we would have been able to by incentivizing direct in-person exchanges.

---

[20]The results are shown in Appendix Figure A.2.b.

4.6 *Diffusion of Information about Salary Versus Seniority*

In this section, we compare the results between the salary and seniority surveys. The key comparisons are presented in Table III. This table presents a series of regressions of some outcomes of interest on a dummy variable that equals 1 if the survey type is seniority and 0 if the survey type is salary.

We start by comparing the misperceptions about peer average salary and peer average seniority. One challenge is that guessing the average seniority of five peers is significantly harder than guessing the average salary of five peers, because there is more variation within peer groups in seniority than in salary. To make the misperceptions more comparable between salary and seniority, we divide the salary (seniority) misperceptions by the within-group standard deviation in salary (seniority).

Columns (1) and (2) correspond to the comparison of misperceptions, using the normalization discussed above. The dependent variable in column (1) corresponds to the average absolute misperceptions, which is 0.707 for salary and 0.469 for seniority, with a p-value of the difference of <0.001. In other words, on average, individuals miss the mark by 0.469 standard deviations when guessing peer seniority but they miss the mark by even more, 0.707 standard deviations, when guessing peer salary.

As a robustness check, the dependent variable in column (2) corresponds to the misperceptions that (hypothetically) would have arisen if the responded reported their own salary (or seniority) as the guess. The results from column (1) and column (2) suggest that the individuals are as good at guessing peer salaries in actuality (average "actual" misperceptions of 0.694, from column (1)) as they would have been if the only information they had was their own salary (average "naive" misperceptions of 0.707, from column (2)). On the contrary, individuals guess their peers' seniority in actuality (average "actual" misperceptions of 0.469, from column (1)) with greater accuracy than they would if the only information they has was their own seniority (average "naive" misperceptions of 0.896, from column (2)). 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 the costs of information acquisition (i.e., information frictions) are higher for salary than for seniority. An alternative interpretation could be that the benefits from information acquisition are 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 benefits for seniority, we should observe a higher willingness to pay for information about seniority than about salary. Column (3) of Table III tests that hypothesis. We find that, on average, there is a large difference, but in the opposite direc-

tion: 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=0.026). This evidence suggests that the differences in misperceptions between salary and seniority must be due, at least partially, to differences in the costs of information search.

Although we cannot rule out other explanations, our preferred interpretation is that employees find it easier to learn about seniority than it is to learn about salaries because of differences in privacy norms.[21] We present three pieces of suggestive evidence on this regard.

The first piece of evidence relies on the subjective data, presented in Figure 8. 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 8 (p-value<0.001 for both *Unacceptable* and *Uncomfortable*).[22]

The second piece of evidence is based on the difference in willingness to search for information. 9.a shows that, when offered an extra week to search for information, one third of participants expected to be fully informed about salaries by the end of that week. In comparison, Figure 9.b shows that twice as many people, nearly two-thirds of participants, expect to be fully informed about seniority by the end of the week.[23]

The third piece of suggestive evidence is based on the revealed-preference measure on willingness to pay for privacy, reported in column (4) of Table III. There is an economically large difference in this outcome. Individuals need compensation of $68, on average, to allow an email revealing their salary to peers. Individuals need compensation of only $29, on average, to allow an email revealing their seniority. This large difference in demand for privacy, which is statistically significant (p-value<0.001), is consistent with the differences in privacy norms.

The last piece of evidence is based on the social distance between employees. As a first measure of social distance, we measure the time working together as the average time the

---

[21]For example, an alternative explanation could be that the firm discourages employees from discussing salaries but not from discussing seniority.

[22]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. Norms about reciprocity, on the other hand, 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.

[23]On the other hand, nearly one-fifth of participants do not expect to improve their accuracy, despite the insensitive nature of the information. This indicates that there are other information frictions besides privacy norms, such as attention and time costs.

subject has worked on the same team with any of the five peers during their past at the bank. The average overlap is 2.8 years. We find that the search frictions associated with asking about seniority decline with the time spent working together, but remain persistently large when asking about salary. In Figure 10 we see a statistically increasing degree of accuracy in the guess, and higher search intensity, as overlap increases; a standard deviation increase in the share of overlapping time together is associated with a decline in misperception of 3.3 percentage points. But this correlation exits only in the case of seniority. By contrast, we see that overlap does not predict accuracy or search when the information exchanged is salary. As a second measure of social distance, we use the Eigenvector centrality of an individual's position in the (bi-directional) graph of emails. Models and empirical tests of social learning have demonstrated that central individuals in a network are best informed (Alatas et al., 2016; Banerjee et al., 2013). We find that higher network centrality is indeed predictive of being more informed about the seniority of peers, but network centrality is uncorrelated with misperceptions about salary. A one standard deviation increase in network centrality predicts a decline in misperception about seniority of 3.4 percentage points, and no change in misperception about peer salary. We interpret this as additional evidence that, consistent with the salary taboo, there are higher barriers to information diffusion about salary than about seniority.

4.7 *Gender Differences*

There is a widespread belief that pay secrecy tends to hurt women and minorities disproportionately, allegedly because women are less likely to search for salary information (Babcock and Laschever, 2009). These claims, however, are mostly based on survey data and anecdotal evidence. 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. However, there is no evidence on whether those survey claims are backed by actual differences in knowledge and information acquisition.[24]

Table IV 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 IV suggests

---

[24]There is evidence of gender differences in the diffusion of other forms of information besides salaries. 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).

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 = 0.217). Column (10) indicates that, according to the incentive-compatible measure, the perceived probability of winning the game is 50% for men versus 39% for women (p-value of the difference = 0.0172). However, the comparison of actual accuracy indicates that these differences in perceived accuracy are misleading. Column (8) of Table IV 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 26.8% versus 34.2% for women (p-value of the difference = 0.168).

Besides this difference in confidence, there are no significant gender differences in any of the other outcomes reported in Table IV. 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.

5. CONCLUSIONS

Barriers to the diffusion of 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 are asked about seniority instead of salary.

Several policies have been enacted to increase pay transparency. Our evidence suggests that some of these policies may be highly valued by employees. 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 employees 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.

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 salary 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% reported that they 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 California, Florida, New York, and other U.S. states, 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, perhaps 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*, 1–9.

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

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

Allcott, H. and J. B. Kessler (2019). The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons. *American Economic Journal: Applied Economics 11*(1), 236–276.

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

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.

Baker, M., Y. Halberstam, K. Kroft, A. Mas, and D. Messacar (2019). Pay Transparency and the Gender Gap. *NBER Working Paper No. 25834*.

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

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*.

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.

Cason, T. N. and C. R. Plott (2014, 12). Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy 122*(6), 1235–1270.

Chandrasekhar, A. G., B. Golub, and H. Yang (2018). Signaling, Shame, and Silence in Social Learning. *NBER Working Paper No. 25169*.

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

Conlon, J., L. Pilossoph, M. Wiswall, and B. Zafar (2018). Labor Market Search With Imperfect Information and Learning.

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 in a simple labor market. *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 r. 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*.

Gee, K. (2017). Ask Me How Much Money I Make. *The Wall Street Journal* (October 26 2017).

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 Womenâs Policy Research*.

Horowitz, J. K. (2006). The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Economics Letters 93*(1), 6–11.

Karni, E. and Z. Safra (1987). "Preference Reversal" and the Observability of Preferences by Experimental Methods. *Econometrica 55*(3), 675.

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. (2019). The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment. *NBER Working Paper No. 25622*.

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, and T. Von Wachter (2019, 2). Firming Up Inequality. *Quarterly Journal of Economics 134*(1), 1–50.

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*.

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

Figure 1: Conceptual Framework

a. Basic Framework

b. Willingness to Search

c. Willingness to Search, by Reward

d. Willingness to Search, Alternatives

e. Willingness to Pay for Information

f. WTP for Info, By Rewards
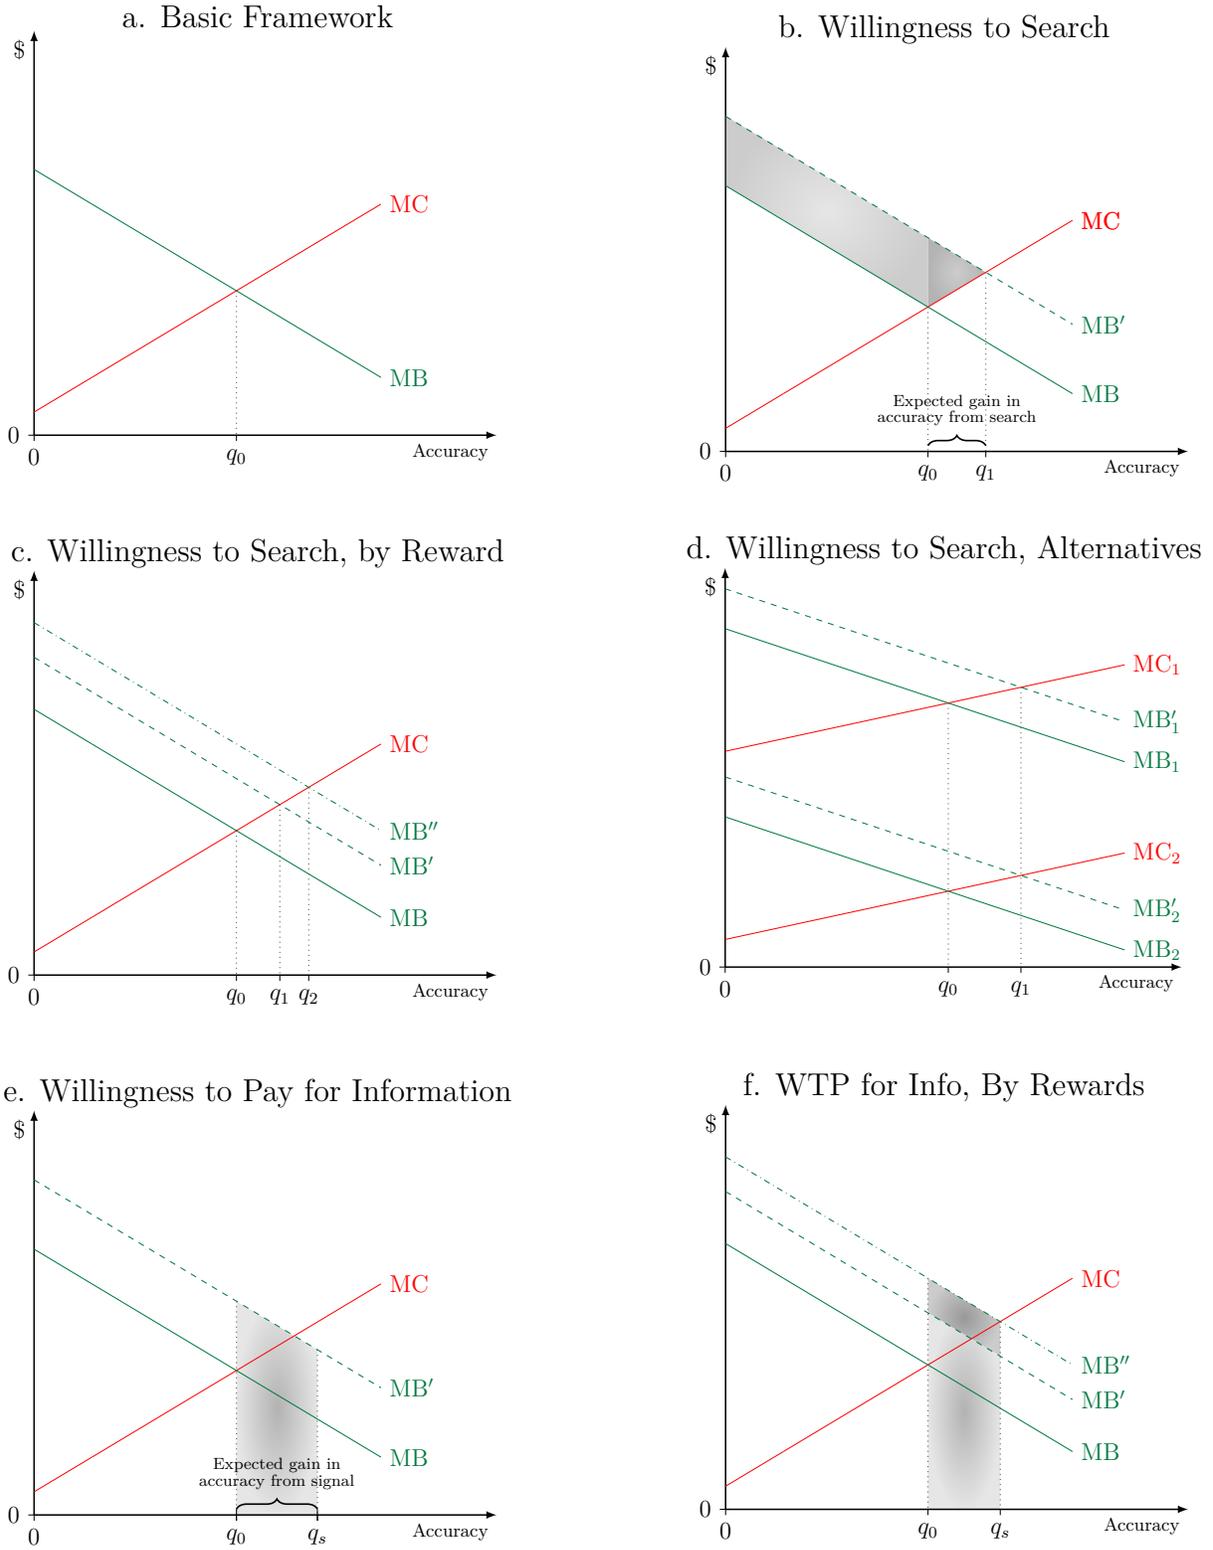
Figure 2: Salary Information: Misperceptions

a. Guess Vs. Own-Salary Benchmark

b. Guess Vs. Peer-Group Benchmark
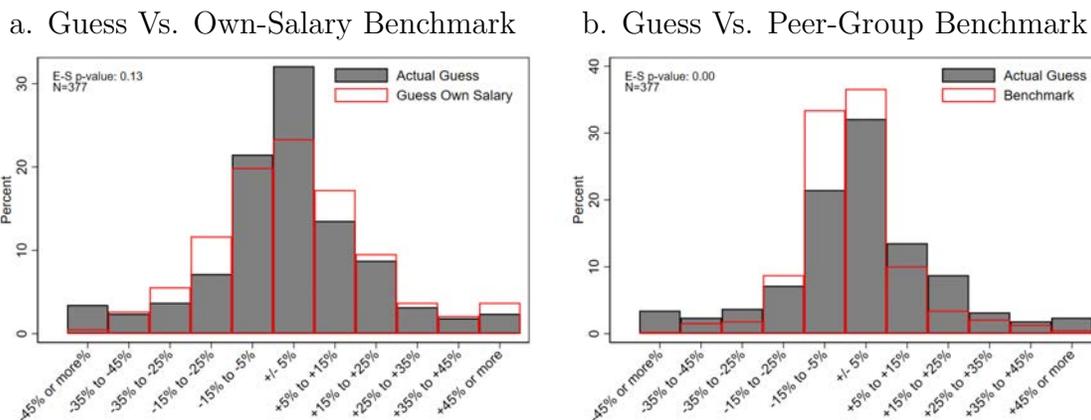
Figure 3: Salary Information: Perceived Accuracy with and without the Extra Week

a. Self-Reported

b. Incentive-Compatible
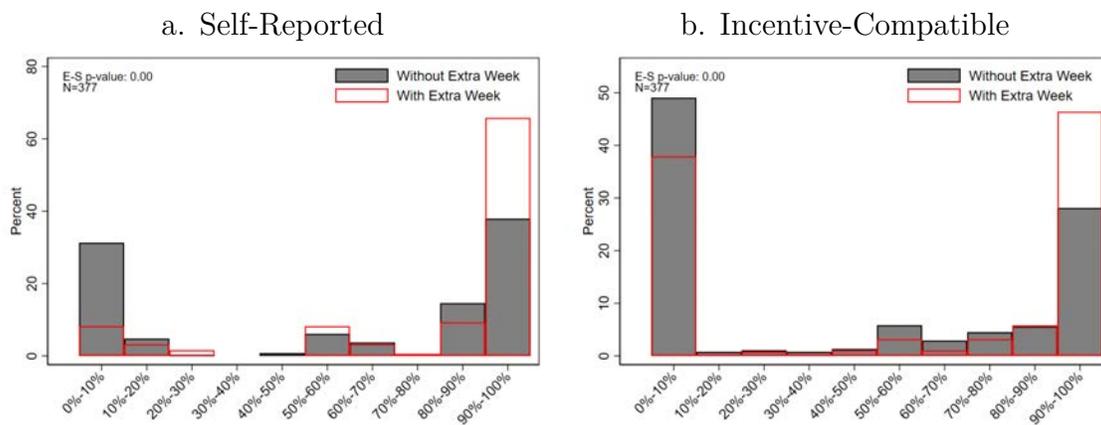
Notes: Panel (a) shows a histogram of the respondent's perceived probability of winning the guessing game, without an extra week to search for information (grey bars) and with the extra week (red bars). Panel (b) is equivalent to the first panel, only that instead of using self-reported probabilities, 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: Salary Information

b. Rational Inattention Test
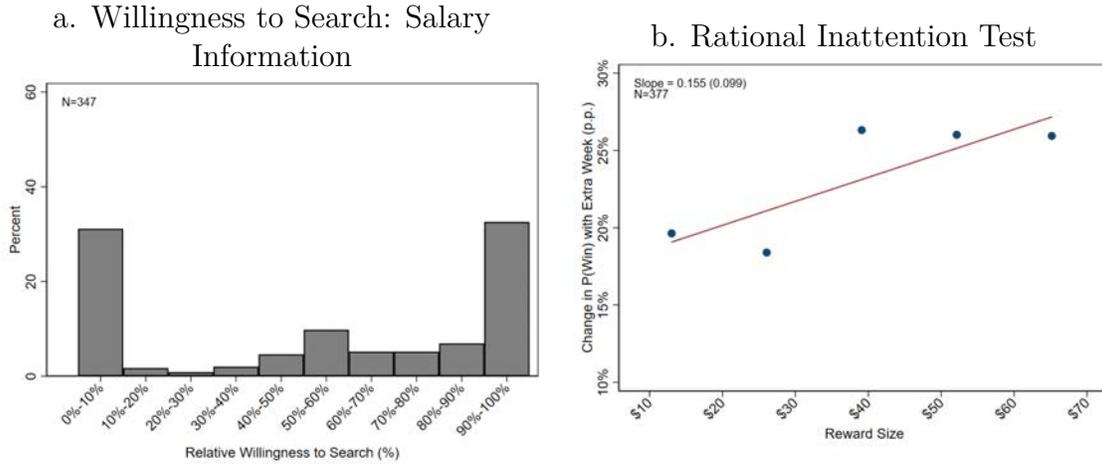
Notes: Panel (a) shows the difference between the probability of winning the game with and without the extra week, divided by one minus the probability without the extra week (this sample excludes 30 individuals who were 100% confident in their initial guess). Panel (b) provides a binned scatterplot 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
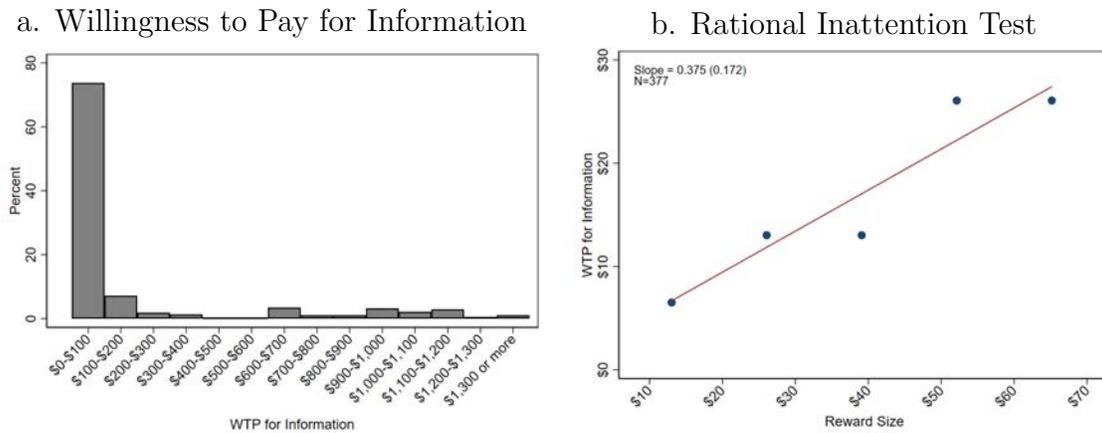
a. Willingness to Pay for Information

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
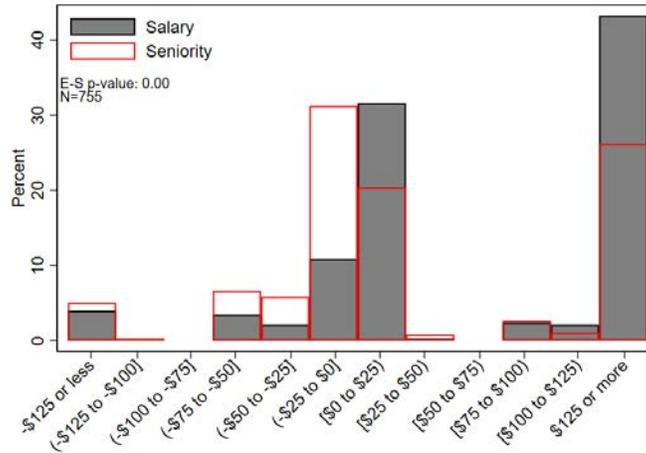
Notes: This histogram shows the distribution of the willingness to pay for privacy: negative values denote the amount the individual is willing to pay to reveal her information to peers, while positive values denote the compensation the individual is willing to give up in order to conceal this information. 15% of participants selected $0, and all of these individuals also indicated that they would not like the email to be sent, so they are included in the bracket [$0,$25).

Figure 7: Willingness to Pay for Privacy by Perceived Relative Standing
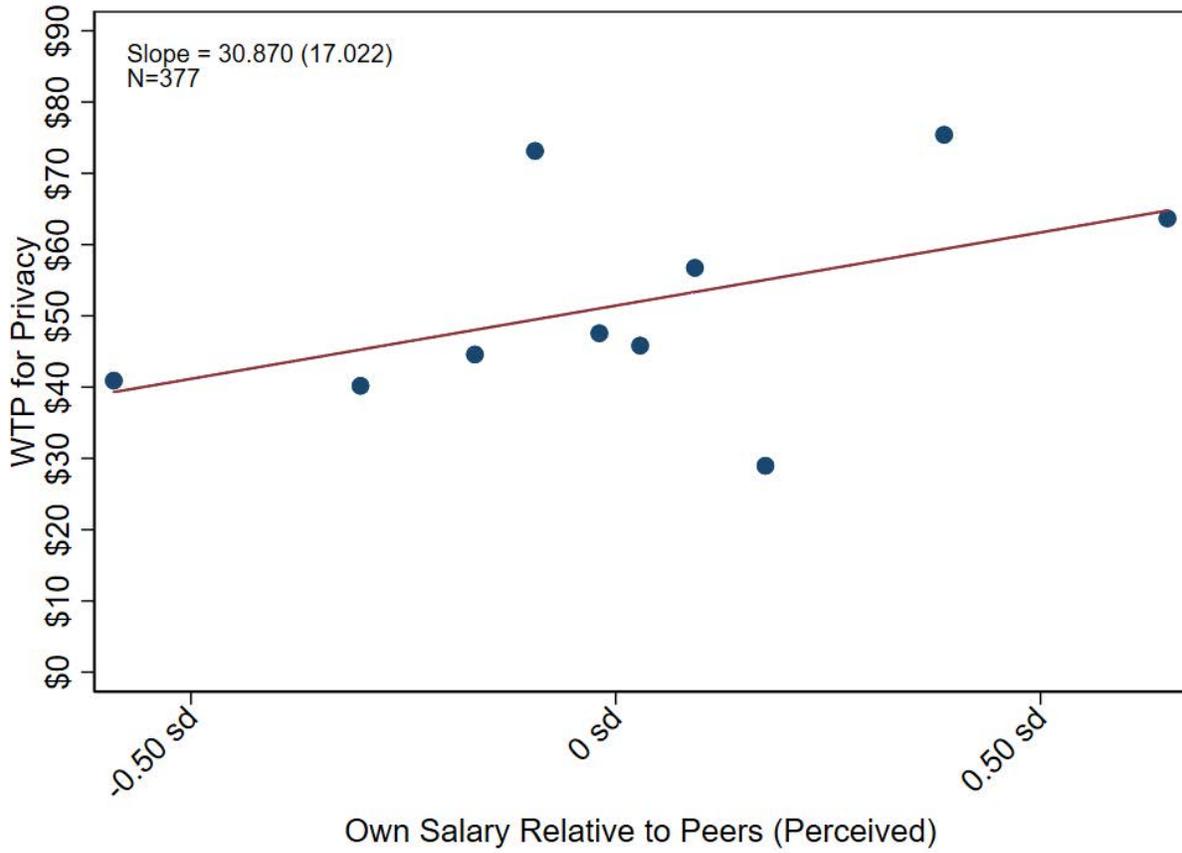
Notes: We provide a binned scatterplot with the relationship between the willingness to pay for privacy and the respondent's perceived relative salary with respect to the reference peer group. Distance from the reference group has been normalized by a standard deviation among peers, and winsorized at the 5th/95th percentiles. The slope is calculated using interval regression with robust standard errors.

Figure 8: Privacy Norms: Salary vs. Seniority

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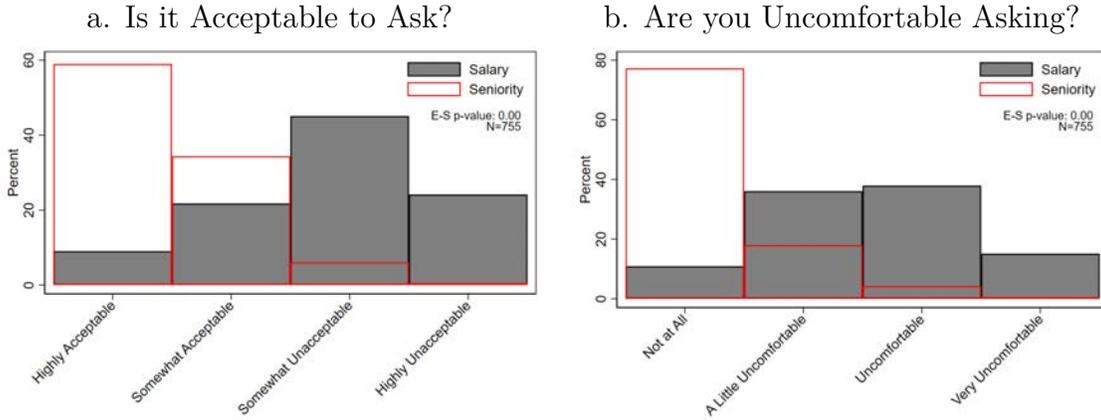
Notes: Panel (a) shows the distribution of responses to the question *Unacceptable*, asking whether it is "socially acceptable to ask someone about their salary/seniority". Panel (b) shows the distribution of responses to the question *Uncomfortable*, eliciting whether the respondent finds it "uncomfortable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

Figure 9: Willingness to Search for Information in the Wild

a. Salary Information   b. Seniority Information
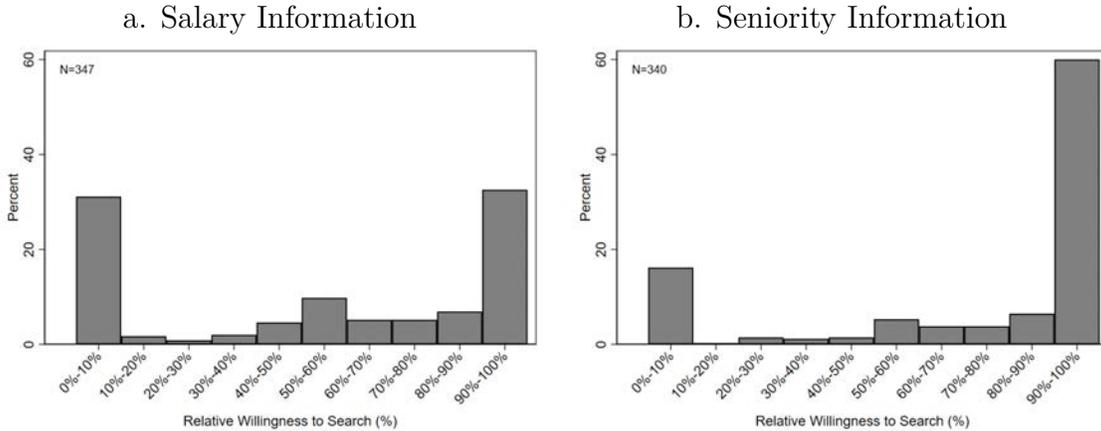
Notes: Panel (a) shows, for those asked about salary information, the difference between the probability of winning the game with and without the extra week, divided by one minus the probability without the extra week (this sample excludes 30 individuals who were 100% confident in their initial guess). Panel (b) shows the same outcome for those asked about the seniority of their peers (excluding the 38 individuals who were 100% confident in their initial guess).

Figure 10: Salary Perceptions and Peer Connectivity

a. Misperceptions with Peer Overlap   b. Misperceptions with Peer Centrality
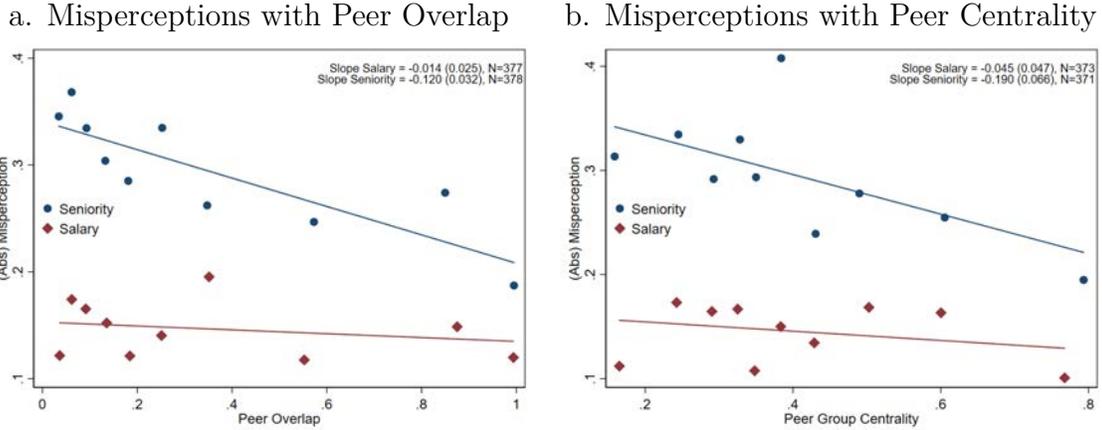
Notes: In panel (a) the binscatter plot and corresponding OLS regression estimates show the relationship between absolute misperception about the average seniority or salary among the five peers (y-axis) and the amount of time that the survey respondent overlapped with the five peers (x-axis). Peer overlap is measured as the average share of the 8 year panel that the survey respondent overlapped at the bank with the five selected peers. We include fixed effects for the unit that the peers and respondent share at the time of the survey. The difference in slopes when jointly estimated is 0.114 (0.0391), p-value= 0.004 In Panel (b) the specification is identical but the dependent variable (x-axis) is the individual's centrality in their peer group measured using the eigenvector centrality of the bi-directional work email graph. The difference in slopes when jointly estimated is 0.145 (0.0842), p-value= 0.085

Table I: 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	0.71	0.74	0.50	0.70	0.72	0.73	0.72	0.75	0.89
	(0.02)	(0.02)	(0.02)		(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	
Age (Years)	29.24	29.39	29.09	0.40	28.80	28.93	29.64	29.23	29.62	0.48
	(0.18)	(0.26)	(0.25)		(0.42)	(0.38)	(0.42)	(0.38)	(0.41)	
College (=1)	0.86	0.86	0.85	0.76	0.83	0.86	0.85	0.89	0.84	0.55
	(0.01)	(0.02)	(0.02)		(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	
Seniority (Years)	4.21	4.29	4.13	0.55	3.82	4.02	4.25	4.45	4.44	0.51
	(0.13)	(0.19)	(0.19)		(0.30)	(0.30)	(0.33)	(0.28)	(0.30)	
Own Salary (Masked)	1.00	0.99	1.01	0.64	0.92	1.00	0.96	1.04	1.06	0.38
	(0.02)	(0.03)	(0.04)		(0.04)	(0.05)	(0.05)	(0.05)	(0.06)	
Observations	755	377	378		135	161	131	167	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 II: Salary Valuations: Heterogeneity by Months Until Negotiation

	(1) Prob. Promotion w/in 3 months	(2) Prob. Raise w/in 3 months	(3) Prob. Family Leave w/in 3 months
WTP>$100	0.086**	0.067*	-0.013*
	(0.044)	(0.037)	(0.007)
Constant	-0.659***	0.095	0.049*
	(0.092)	(0.070)	(0.028)
Mean Dep. Var. (WTP<$100)	0.14	0.09	0.01
Std. Dep. Var.	0.36	0.31	0.09
Observations	370	367	374
R^2	0.199	0.010	0.010

Notes: All linear probability models are estimated using OLS. Significant at *10%, **5%, ***1%. Standard errors in parentheses clustered at the position level. The dependent variable in Column (1) and Column (2) is equal to 1 if the respondent received a promotion or raise within the 3 months following the survey, and 0 otherwise. The dependent variable in Column (3) is equal to 1 if the person took extended family leave within 3 months of the experiment, and 0 otherwise. We control for the salary of the respondent. We exclude employees who quit within 10 months following survey completion.

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

	Abs. Error (Std.)		WTP ($)	
	Actual (1)	Naive (2)	Signal (3)	Privacy (4)
Seniority Dummy (=1)	-0.238***	0.202***	-49.120**	-38.931***
	(0.056)	(0.052)	(21.894)	(7.636)
Constant	0.707***	0.694***	178.641***	67.538***
	(0.049)	(0.034)	(18.454)	(5.768)

Notes: N=755. 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 side 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 IV: Salary Information: Heterogeneity by Gender

	Attitudes			Error		Absolute Error	
	(1) Uncomf.	(2) Unacc.	(3) Recipr.	(4) Actual	(5) Naive	(6) Actual	(7) Naive
Female (=1)	-0.063	0.141	-0.000	0.028	-0.079***	0.003	-0.042**
	(0.099)	(0.101)	(0.033)	(0.022)	(0.030)	(0.017)	(0.021)
Constant	1.472***	1.056***	0.889***	-0.022	0.064**	0.145***	0.190***
	(0.083)	(0.082)	(0.027)	(0.019)	(0.027)	(0.015)	(0.020)

	Accuracy (pp)	Perceived Acc. (pp)		Extra Week (pp)		WTP ($)	
	(8) Actual	(9) Direct	(10) Indirect	(11) Direct	(12) Indirect	(13) Signal	(7) Privacy
Female (=1)	0.073	-0.057	-0.107**	-0.025	0.017	19.369	-16.765
	(0.050)	(0.043)	(0.052)	(0.028)	(0.033)	(35.612)	(12.778)
Constant	0.269***	0.598***	0.504***	0.603***	0.247***	164.821***	81.599***
	(0.046)	(0.039)	(0.043)	(0.039)	(0.037)	(30.623)	(11.794)

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 I. Column (5) corresponds to the final sample of 755 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 29.2 vs. 30.5 years old, the share of college graduates is 86% vs. 85%, and the mean seniority is 4.2 vs. 4.8 years.

The only noticeable difference between columns (1) and (5) of Table I 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.

Figure A.1: Willingness to Pay for Salary Information by Elicitation Method

a. Open Bidding BDM | b. Multiple Price-List Menu BDM
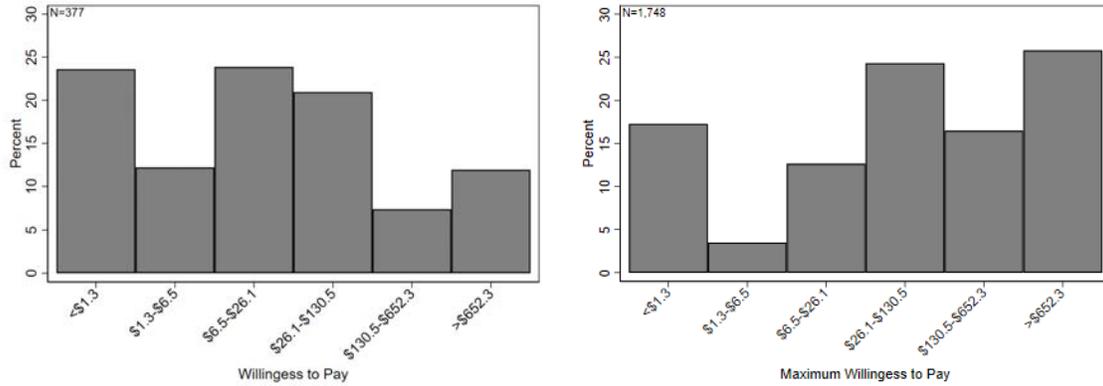
Notes: Panel (a) shows the distribution of the willingness to pay for information about the average salary among a sample of five peers, as measured by the respondent's incentive-compatible bid using the open bidding BDM method. Panel (b) shows the distribution of the willingness to pay for information about the average salary among peers, using the multiple price-list menu BDM method. The sample is restricted to the subset of respondents with consistent responses across the five price scenarios. Study participants are a non-overlapping representative sample from the same institution. The graph is replicated from Cullen and Perez-Truglia (2018).

Figure A.2: Willingness to Pay for Privacy by Perceived Relative Standing
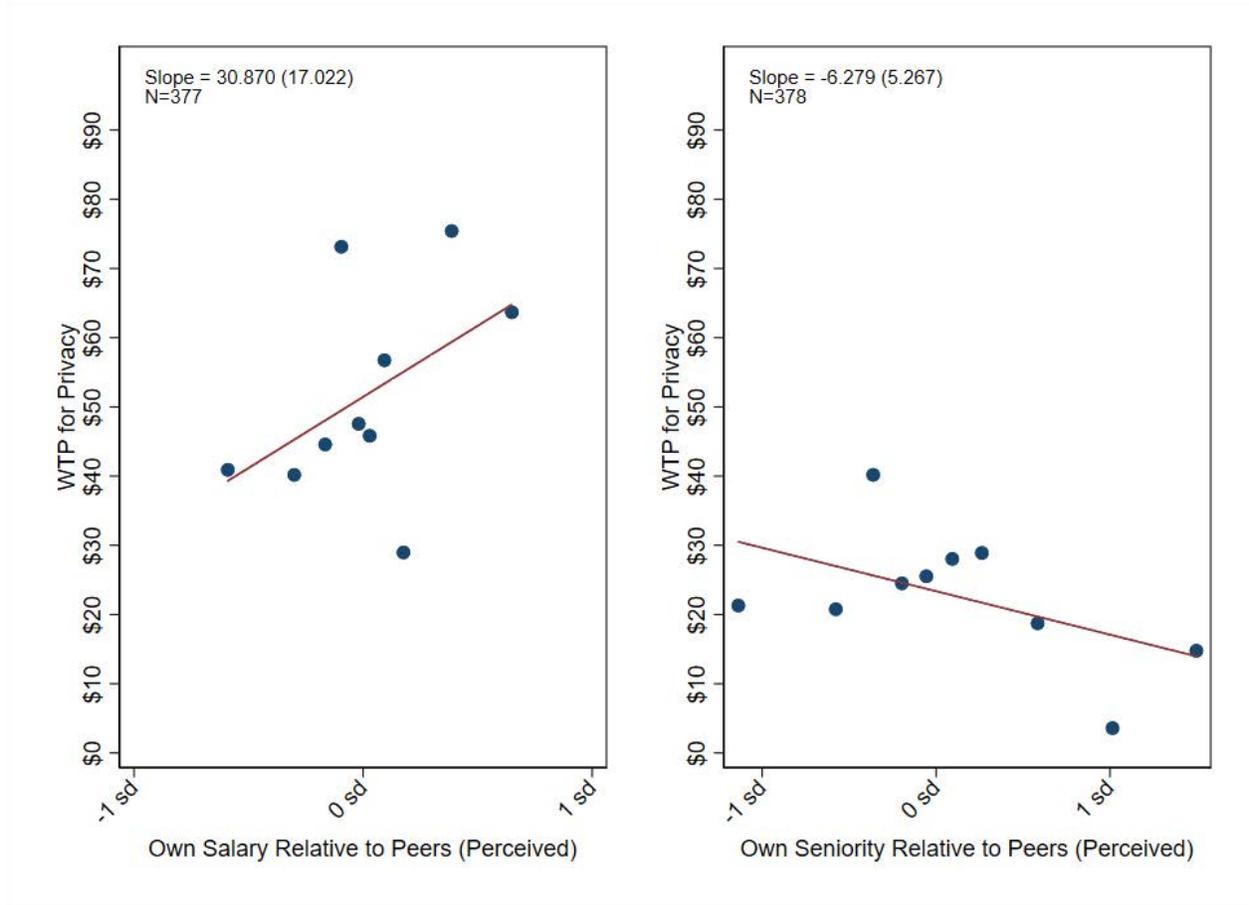
Notes: Panel (a) shows a binned scatterplot with the relationship between the willingness to pay for privacy and the respondent's perceived relative salary with respect to the reference peer group. Distance from the reference group has been normalized by a standard deviation among peers, and winsorized at the 5th/95th percentiles. The slope is calculated using interval regression with robust standard errors. Panel (b) shows the same binscatter for individuals asked about seniority.

Table I: 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)	30.49	30.90	29.18	29.15	29.24
	(0.06)	(0.07)	(0.11)	(0.14)	(0.18)
College (=1)	0.85	0.85	0.85	0.85	0.86
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Seniority (Years)	4.75	4.85	4.44	4.59	4.21
	(0.04)	(0.05)	(0.09)	(0.11)	(0.13)
Own Salary (Masked)	1.42	1.54	0.99	0.99	1.00
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Observations	(Masked)	(Masked)	1,899	1,144	755

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. To acknowledge your help, you will be eligible to monetary rewards, for a minimum of $10.

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

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]

The survey rewards will be deposited automatically in your bank account within 2 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 ([Information Technology's Extension Number]).

You were selected at random to receive this invitation, and all your 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]

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 Harvard University and the University of California at Los Angeles. This survey will teach us about how [Bank name]'s employees learn about their workplace, earnings, and career prospects.

As a reward, you will receive at least $10. In addition to this payment, you will qualify to earn additional rewards between $0 and $100.

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 $

>>

0 2 5 9

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 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/Tenure] among the 5 selected peers as of December 2017?

0   [$/Years]

>>

You have the opportunity to buy the following information: the average salary 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 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 | $ |

>>

X

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/Tenure] 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.

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/Tenure] to some of your peers.

We could send an email to your 5 selected peers revealing your [Salary/Tenure], 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 to your peers?

○ Not at all

○ A little uncomfortable

○ Uncomfortable

○ Very uncomfortable

If you ask a peer about his or her salary, would you expect this peer to ask you about your salary?

○ Yes

○ No

Is it socially acceptable to ask someone about their salary?

○ Highly unacceptable

○ Somewhat unacceptable

○ Somewhat acceptable

○ Highly acceptable

>>