



Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpubeThe salary taboo privacy norms and the diffusion of information [☆]Zoë Cullen ^a, Ricardo Perez-Truglia ^{b,*}^a Harvard University, Rock Center 211, Boston, MA 02163, United States^b University of California, 545 Student Services Building #1900, Berkeley, CA 94720, United States

ARTICLE INFO

Article history:

Received 26 May 2022

Revised 31 March 2023

Accepted 9 April 2023

Keyword:

Information diffusion

Salary

Privacy

Inequality

Transparency

Gender

ABSTRACT

The diffusion of salary information has implications for labor markets, such as wage discrimination policies and collective bargaining. Access to salary information is believed to be limited and unequal, but there is little direct evidence on the sources of these information frictions. Social scientists have long conjectured that privacy norms around salary (i.e., the “salary taboo”) play an important role. We provide unique evidence of this phenomenon based on a field experiment with 755 employees at a large commercial bank in Southeast Asia. We show that many of its employees are both unwilling to reveal their salaries to coworkers and reluctant to ask coworkers about their salaries. These frictions persist, in smaller magnitude, when sharing less sensitive information on seniority. We discuss implications for pay transparency policies and the gender wage gap.

© 2023 Elsevier B.V. All rights reserved.

You may know about your colleague’s sex life, your friend’s drinking problem, or what your neighbor really thinks of her mother-in-law. But you probably don’t know what they take home in each paycheck. — Margaret Littman, *Working Woman*, 2001.

1. Introduction

Employees use salary information in several ways, such as to negotiate raises or compare positions or employers. As most employers provide limited salary information, employees’ knowledge about it depends largely on their ability to communicate with each other. Recent evidence suggests that the diffusion of salary information can be highly imperfect (Cullen and Perez-Truglia, 2022). Social scientists claim that these information frictions may

stem from a “salary taboo,” which is a social norm around salary privacy that discourages coworkers from revealing or asking others about salary information. The source of this privacy norm is unclear. Some researchers conjecture that individuals are uncomfortable discussing pay because they equate it with discussing self-worth (Trachtman, 1999; Edwards, 2005). It is also possible that highly paid individuals may want to conceal their salary to avoid embarrassment, resentment, retaliation, people asking to borrow money, or expectations that they should work harder. Despite the potential importance of privacy norms for labor markets, there is little *direct* evidence on their role in the diffusion of salary information. In this study, we fill this gap by conducting a field experiment with a multibillion-dollar corporation.

The diffusion of salary information, or lack thereof, has broad labor market implications. For example, information frictions can facilitate workplace discrimination, increase employers’ market power (Danziger and Katz, 1997; Cullen and Pakzad-Hurson, 2021), and hinder collective bargaining and unionization (Corbett, 2002).¹ Understanding the precise sources of these information frictions also is relevant for ongoing policy debates. For example, current U.S. legislation intended to promote pay transparency punishes employers that retaliate against employees who discuss their salaries (Siniscalco et al., 2017; National Labor Relations Board, 2022). This policy can be effective, assuming that the reason employees do not discuss salaries is that they fear being

^{*} We thank the editor (Robert Metcalfe) and three anonymous referees for their excellent feedback. We are thankful for numerous comments from colleagues and seminar discussants. In particular we would like to thank Emily Breza, Kai Barron, Lauren Cohen, Mu-Jeung Yang, Irene Lo, Rediet Abebe, Bobby Pakzad-Hurson, Shane Greenstein, Teresa Amabile, Pietro Ortoleva, Edward Lazear, Chris Stanton, Robert Gibbons, David Slusky, Sera Linardi, Olga Stoddard, Teodora Boneva, Silvia Saccardo, and Clayton Featherstone. Katherine Fang, Anh Nguyen, Dylan Balla-Elliott and Romina Quagliotti provided excellent research assistance. This project was reviewed and approved by the Institutional Review Board at University of California Los Angeles (IRB#18-001467).

^{*} Corresponding author.

E-mail addresses: zcullen@hbs.edu (Z. Cullen), ricardotruglia@berkeley.edu (R. Perez-Truglia).

¹ For a discussion on wage inequality, collective bargaining, and unionization, see Card et al., 2004.

caught by their employers (Gely and Bierman, 2003; Hegewisch et al., 2011). However, if employees are unwilling to discuss salaries due to privacy norms, they may fail to learn information regardless of whether they are discouraged by their employers. If the salary taboo exists, then other policies would be needed to guarantee that employees have access to salary information.

Most employees around the world report a desire to be better informed about the salaries of their coworkers, but they also report that they rarely discuss salaries with each other (Glassdoor, 2016; PayScale, 2018). Although these patterns are indicative of information frictions, they are not necessarily due to a salary taboo. For example, employees may be uninformed because the cost of acquiring salary information, in terms of time and energy spent, outweighs the expected benefits. To explore this issue, we partnered with a large firm to conduct a field experiment with its employees. We designed the experiment to disentangle the role of privacy norms in how employees share their own salary information and inquire about their coworkers' salaries.

We assess the salary taboo directly using measures of social norms around privacy. We ask employees whether it is socially acceptable to ask coworkers about their salaries and whether they feel comfortable doing so. To address the usual concerns with subjective data, we designed a revealed-preference measure of preferences for privacy. Measuring these preferences in a natural, high-stakes context presents challenges, such as verifying whether the salary information actually was shared. To address this challenge, we offer subjects the opportunity to reveal their own salary to a specific set of five peer coworkers, listed by first and last names. We use the standard definition of peer coworker: another employee who shares the same position title and works in the same organizational unit (Card et al., 2012; Cullen and Perez-Truglia, 2022). For example, a junior researcher from investment banking is offered the opportunity to share his or her salary information with five other junior researchers working in investment banking. The information can be shared effortlessly over email, and we offer rewards in exchange for revealing (or concealing) their salary information. This design provides an incentive-compatible elicitation of willingness to share salary information.

Additionally, we measure whether employees are willing to search for information about the salaries of coworkers. This elicitation presents several challenges of its own. For example, we want to avoid concerns about revealing information sources to the experimenter or about repercussions from management for searching for information. To address those challenges, we require employees to guess the average salary of a random sample of five peer coworkers, listed by first and last names, as part of a game endorsed by the firm. If the subject's guess falls within 5% of the truth, he or she receives a monetary reward from \$13 to \$65.²

Next, we offer respondents the opportunity to have an extra week to search for salary information on their own. This additional time would allow them to improve their guesses and thus increase the chances of winning the guessing game. For instance, subjects could use the extra time to reach out to the five peers on the list and ask about their salaries. We elicit the probabilities of winning the guessing game with and without the additional week, using both non-incentivized and incentivized methods. The additional probability of winning with the extra week should measure the employees' willingness to search for information. For example, if the probability of winning the game increases from 50% to 100%, that would indicate that the individual expects to acquire full information during the extra week. On the contrary, if the proba-

bility remains at 50%, that would indicate that the individual does not expect to find any new information.

To further assess the role of privacy norms, we randomize subjects into one of two survey types: the salary survey or the seniority survey. The surveys are identical except that for every question in the salary survey, the seniority survey replaces the word "salary" with the word "seniority". For example, the salary survey asks employees whether it is socially acceptable to ask peers about their salaries, and the seniority survey asks whether it is socially acceptable to ask peers about their seniority.

Like salary information, seniority information can be useful in making important career choices. For example, anecdotal evidence suggests that employees in this firm use their relative seniority to assess whether they are due for a promotion. Just like when acquiring information about salary, acquiring information about seniority involves investing time and effort. We conjecture that privacy norms around seniority will be much weaker than privacy norms around salary, however.³ If there are no privacy norms around seniority information, then the comparison between salary and seniority information could reflect the effect of privacy norms.

We conduct the experiment with a sample of 755 employees from a large commercial banking firm in Southeast Asia (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 is arguably relevant to study information frictions: it does not disclose salary information, its employees reportedly desire more salary transparency, and its employees rarely discuss salaries with their coworkers. Several studies show that these institutional factors 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).

We find that, consistent with the salary taboo hypothesis, employees find salary discussions to be a sensitive topic: 69% find it socially unacceptable to ask coworkers about their salary, and 89% would feel uncomfortable if they had to ask a coworker about their salary. By comparison, employees do not find discussion about seniority to be nearly as sensitive as discussions about salary: only 7% find it unacceptable to ask coworkers about their seniority, and only 23% would feel uncomfortable asking a coworker about his or her seniority.

The revealed-preference measures are consistent with the subjective measures of privacy concerns. Regarding the willingness to share their own information, we find that compared to a minority of employees (20%) who prefer to share personal salary information with five of their peers, most (80%) prefer to conceal this information. Moreover, those who prefer not to reveal their salaries have high willingness to pay for their privacy. For example, 40% of respondents would not reveal their salary information for the maximum incentive in the experiment, \$125, equivalent to roughly four days' worth of pay. In contrast, we find a significantly higher willingness to share seniority information, with only 23% of employees refusing to reveal their seniority in exchange for \$125.

Although privacy norms seem to play a significant role, they are probably not the only factor preventing individuals from sharing their salary information. In particular, some preferences for secrecy may be strategic. For example, consider an employee who is paid more than their peers and who is deciding whether to reveal his or her salary to peers. On the one hand, revealing the information may gain respect and admiration from peers. On the other hand, it may generate resentment or even undermine the employee's leverage in future salary negotiations. Indeed, we find evidence

² These amounts, as well as all other monetary amounts discussed in the paper, are converted to U.S. dollars using PPP-adjusted exchange rates from February 2018.

³ For example, relative salary is believed to be sensitive, because it reflects one's worth (Gely and Bierman, 2003; Hegewisch et al., 2011). Relative seniority, however, may not be as informative about one's worth.

for the second channel: those who perceive themselves to be higher relative earners are less willing to reveal their salaries to peers.

If privacy norms create information frictions, we would expect employees to be more informed about seniority than about salary. To measure misperceptions, we compare the employees' guesses about the average salary of peers to the true averages from the administrative records of the firm. Guesses about the average salary of the five peers have a mean absolute percentage error of 14.6%. This level of misperception is what we would expect if employees had used their own salary to guess the salaries of their peers, thus suggesting that employees have little information about anyone's salaries but their own. In contrast, employees are substantially more accurate when guessing seniority than they would be if they just reported their own seniority. One possible interpretation is that employees have access to other information about seniority but no access to other information about salaries other than their own salary. While speculative, this constitutes suggestive evidence that there is less social learning for more sensitive information.

Next, we turn to the willingness to search for information. We find that most individuals are willing to search for seniority information when offered financial incentives to do so. Two-thirds of participants who express uncertainty in their initial guess about seniority expect to return after a week with the correct answer. However, willingness to search falls significantly when subjects are asked about (more sensitive) salary information, with less than one-third expecting to return after a week with an answer within 5% of the true average salary. Furthermore, 25% of respondents are not willing to search for salary information on their own, even when offered the highest prize for accuracy (\$65).

We design a validation test to measure willingness to search for information. We randomize the reward of the guessing game to take one of five different values ranging from \$13 to \$65. According to models of rational information acquisition, we expect higher rewards to increase employees' willingness to search for information, because they stand to gain more from it (Woodford, 2001; Sims, 2003; Mankiw and Reis, 2002; Reis, 2006). Consistent with this prediction, we find that employees have significantly higher willingness to search when they are randomly assigned to higher reward amounts.

Prior theoretical and empirical studies suggest that, in the presence of social learning, employees who are more central in the network (or better connected to the specific five peers they have to guess about) should have fewer misperceptions (Banerjee et al., 2013; Mobius and Rosenblat, 2014; Alatas et al., 2016). We show that employees who are more central (or better connected) do have fewer misperceptions about peers' seniority, but they do not have fewer misperceptions about peers' salaries. This evidence again suggests that the social diffusion of less sensitive information (seniority) is greater than the social diffusion of more sensitive information (salary).

Our last result relates to the widespread belief that pay secrecy and privacy norms disproportionately hurt women (Babcock and Laschever, 2009). For example, survey data indicate that women feel less informed than men about the salaries where they work (Glassdoor, 2016; Cullen and Pakzad-Hurson, 2021). 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 difference 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 are also small, statistically insignificant, and precisely estimated: female and male employees are equally willing to say that they are uncomfortable asking others for their salaries and equally willing

to reveal their own salaries to their peers or to search for information.

Our results are relevant for several recent pay transparency policies that have been enacted or discussed as they suggest that some of these policies may be more useful than others. 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 limited effects on the diffusion of salary information, at least in the short-term, because employees would still be reluctant to inquire about salaries due to privacy norms. Also, our evidence on employees' demand for privacy suggests that employees can be hurt significantly by transparency policies that result in disclosing non-anonymous information. For instance, the salaries of U.S. state employees are made publicly available and easily accessible online. These databases include not only salary information but also the full names of employees. Our findings suggest that this disclosure can have a direct utility cost to the individuals whose information is disclosed.

Our study relates to various strands of literature. For example, we contribute to a small but growing literature showing that increased pay transparency affects employee behavior and satisfaction (Card et al., 2012; Mas, 2016; Mas, 2017; Breza et al., 2018; Perez-Truglia, 2020; Cullen and Pakzad-Hurson, 2021; Baker et al., 2023), including one study using a different experiment at the same firm used in our study (Cullen and Perez-Truglia, 2022). This evidence suggests that individuals have imperfect access to information, but it does not address the sources of these information frictions.⁴ We contribute to this literature by measuring the role of privacy norms. In this sense, this study is also related to the literature on wage discrimination, particularly the widespread view that information frictions hurt minorities disproportionately (Colella et al., 2007; Phillips, 2009). Our evidence does not support this view: we find that women and men face similar frictions and have similar degrees of misperceptions.

Our study also contributes to the literature on 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 (Habyarimana et al., 2007; Mobius and Rosenblat, 2014). Some of these studies artificially create incentives for information diffusion. For instance, Mobius et al., 2015 recruited college students to play a "treasure hunt" game in which they earned rewards by collecting information from peers. Other studies exploit natural incentives for information diffusion. For example, Beaman et al., 2018 seeded useful information about composting and measured its diffusion in an agricultural network. These papers provide evidence that, even in settings where information is mutually beneficial, its diffusion is highly imperfect.

Our contribution to the social learning literature is twofold. First, we provide evidence of a previously undocumented phenomenon of privacy norms hindering social learning. Beyond salary discussions, information taboos may hinder information diffusion in a range of topics. For example, clients might not want to talk to each other about the prices that they paid, which enables price discrimination, and individuals with mental health afflictions may not want to discuss this information with others, which limits social learning about treatments.⁵ Second, we contribute a new

⁴ In Cullen and Perez-Truglia (2022), we document that employees have large misperceptions about the salaries of managers and peers, and when we provide that information for free to a random sample of employees, it does not spread to other employees in the firm. In this study, we follow up on that work by studying the sources of information frictions.

⁵ Relatedly, Chandrasekhar et al., 2019 investigate whether shame hinders information acquisition (i.e., whether individuals avoid seeking helpful information when doing so would signal ignorance).

method for measuring the willingness to search for information from and share information with others. Our method is unique in that it does not require participants to reveal explicitly that they have searched in any incriminating way, nor does it incriminate the people who would have shared the information. As such, our method could be used to test the spread of sensitive information in other settings where it could be dangerous for individuals to reveal their connections, such as politically sensitive scenarios.

Our paper also relates to the literature on the economics of privacy (Acquisti et al., 2016). For example, Goldfarb and Tucker, 2012 show that, even in anonymous internet surveys, some respondents refuse to reveal information about their incomes and demographics. Athey et al., 2017 and Adjerid et al., 2013 study the demand for privacy in the cryptocurrency market and find 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 large heterogeneity in preferences for privacy, with some individuals willing to pay to reveal their salary to peers rather than conceal it.

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

2. Conceptual framework

The decision to seek out career information from colleagues depends on the perceived instrumental value of acquiring the information, such as negotiating a higher pay or career planning. For example, in the case of salary information, if you find out that the salary of your peers is 10% higher than expected, then there is some chance that you will be able to get a 10% raise by bringing this to the attention of a manager. If an employee learns that the salary is less than expected, she may revise expectations of earnings at that firm downwards, and perhaps even seek alternative offers outside the firm. Similarly, in the case of seniority information, learning that the experience of coworkers is higher or lower than anticipated will help to formulate expectations about time to promotion and salary negotiations. The instrumental value of information must be positive in so far as information improves career decisions, akin to the positive value of advice in Chandrasekhar et al. (2019).

On the other hand, career-related information comes at a cost. Individuals must invest limited time and attention in seeking out colleagues and asking them questions. Such costs are present regardless of the type of information sought out. When the information sought is considered sensitive, as we conjecture to be the case for salary information, an additional cost may be incurred that we refer to as the "taboo" cost. Taboo costs may stem from violating a social norm around appropriate matters of conversation (Fershtman et al., 2011). The employer may enforce the taboo either through explicit or implicit shunning of conversation topics. The taboo could also originally stem from the presence of other frictions, including privacy concerns. Once established, however, a taboo can exert force through channels like shaming. Lastly, the taboo costs may not be equal for all employees. Not only are individuals raised in different environments, but the taboo may also differ according to the status of individuals. Premised on the fact that approximately 90% of participants expect that asking information of others will result in being asked the same question in return, we must allow for the possibility that acquiring new salary

information requires revealing one's own information. Indeed, revealing oneself to be high or low status can be associated with distinct taboo costs that are not necessarily symmetric (Bursztyn, 2015).

For a more formal representation of the costs and benefits of acquiring information, Fig. 1 provides a graphical framework. Assume that an employee's utility depends on the accuracy of their information about the career status of peers. Employees can take actions in order to improve the accuracy of their beliefs about peer status. Those actions, however, 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. Fig. 1 provides a graphical representation of this cost-benefit analysis. The x-axis of Fig. 1.a 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.

Fig. 1.a shows the Marginal Benefit (MB) and Marginal Cost (MC) curves. We assume that the marginal benefit curve is downward sloping: as accuracy rises, additional precision benefits the employee less. The MC curve represents the traditional information acquisition costs, like energy, time and attention. The MC curve is upward sloping because gaining accuracy gets more difficult as one approaches full certainty. For example, when the individual has a single piece of information, acquiring another piece of information doubles the accuracy. However, when the individual has dozens of datapoints, acquiring a single datapoint will lead to a minor increase in accuracy, so the individual will need to seek out more information to attain the same gain in accuracy.

The MC_{Taboo} curve represents the costs of acquiring sensitive information. Relative to the MC curve, the MC_{Taboo} curve is shifted upwards, representing the fact that information is more costly: on top of the time and energy required, the individual has to pay an additional cost related to breaking privacy norms. The MC_{Taboo} is also steeper than the MC curve. For example, a given employee may be more comfortable asking some peers than others. Assuming that this employee will start by asking the peer who he or she feels most comfortable asking, then she will move to the second most comfortable, and so on, this will result in an increasing marginal cost of breaking the taboo.

Fig. 1.a shows the accuracy level we would expect to observe in the case of seeking sensitive and non-sensitive career information. We assume for simplicity that signals of both types of information provide the same marginal benefit, but we allow for their marginal costs to diverge.⁶ q_0 shows where the marginal benefit curve from career-related information intersects with the marginal cost curve including only traditional seeking costs. We expect employees to search for information until this point before providing any additional incentives. Furthermore, q_0^T marks the equilibrium level of accuracy with respect to sensitive, or Taboo, information. These stylized curves highlight that the $q_0^T < q_0$, which in our setting predicts that prior beliefs about the salaries of peers will be less accurate than prior beliefs about the seniority of peers. This.

Prediction 1. Guesses about salary information will be less accurate than guesses about seniority information ($q_0^T < q_0$).

⁶ This assumption is indeed for the sake of simplicity, but is also founded on the fact that seniority is an important part of the salary determination formula, and therefore an essential fact to know in order to negotiate by comparing to others, in much the same way that salary information can be useful.

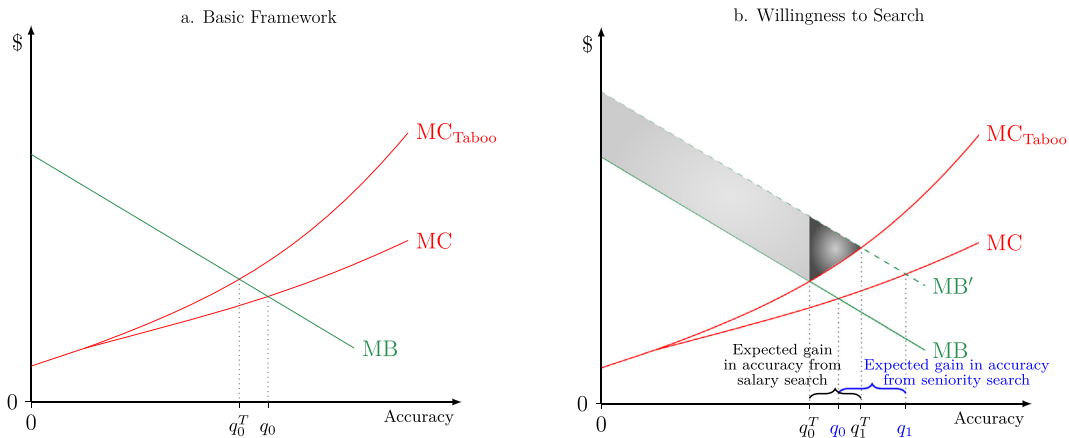


Fig. 1. Conceptual framework.

Fig. 1.b shows how accuracy levels for salary and seniority might shift once we introduce additional rewards for accurately reporting characteristics of peers. The marginal benefit curve shifts upwards and the framework predicts that employees will be willing to carry out additional search. This leads to the following prediction:

Prediction 2. Accuracy of guesses will increase with additional rewards ($q_0^T < q_1^T$ and $q_0 < q_1$).

Employees are now willing to seek out information until where the new marginal benefit curve intersects the marginal cost curve, being q_1 in the case of non-sensitive information and q_1^T in the case of taboo information. In this framework, where taboo costs cause the marginal cost curve to rise steeply above and beyond traditional seeking costs, the additional reward creates a stronger incentive for employees to seek out non-sensitive information as compared to sensitive information. This is summarized in the following prediction:

Prediction 3. Additional rewards create a stronger incentive for employees to search for seniority information as compared to salary information ($(q_1 - q_0) > (q_1^T - q_0^T)$).

We will come back to this framework in the next section, to motivate and illustrate the research design.

3. Research design

Appendix C includes a full sample of the survey instrument. In this section, we describe the main features of the survey. Table 1 provides a summary of the key survey questions, which are explained in detail below.

3.1. Survey types: salary vs. seniority

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, 2022). For instance, we define the peers of a bank teller from one specific branch as the other tellers at that 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, allow-

ances, 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.

To simplify exposition, in the remaining of the section we only describe the questions corresponding to the salary type. 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. 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.

3.2. 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 (e.g., for privacy) 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 (the respondent’s privacy). The rules are as follows. The respondent’s bid is compared to a price that is determined by a

Table 1
Summary of key survey questions.

Goal	Question
Guess	What is the average [Salary/Seniority] among the 5 selected peers?
Perceived accuracy	What do you think is the probability that your guess will fall within +/-5% of the truth?
Willingness to pay for information	What is the maximum amount you would be willing to pay for the information (the average [Salary/Seniority] among a random sample of 5 peers)?
Willingness to search for information	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? [Without the extra week] How much should we pay you so that you give up the right to play this game? [With the extra week] Now, with the additional week, how much should we pay you so that you give up the right to play this game?
Privacy norms	Is it socially acceptable to ask someone about their [Salary/Seniority]? Do you find it uncomfortable to ask information about [Salary/Seniority] to your peers? If you ask a peer about his/her [Salary/Seniority], would you expect this peer to ask you about your [Salary/Seniority]?
Preference for privacy	Would you want us to send an email to your 5 peers revealing your [Salary/Seniority]? [Yes/No] [If "Yes"]: What is the maximum amount you would pay us to send this email? [If "No"]: What is the minimum amount that you would accept to let us send this email?

Notes: The table summarizes the questions in the survey used to measure the main outcomes of interest. A sample of the full survey experiment is attached as Appendix C.

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 make 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 theory, to ensure BDM to be incentive compatible the subjects must know that the probability is positive but it does not matter exactly what the value of the probability is. There is abundant evidence that this is true in practice too: the results from elicitation are similar regardless of whether one uses low or high probabilities, as long as the probabilities are greater than zero (Carson and Groves, 2007; Charness et al., 2016). We executed the BDM choices for 1% of the subjects invited to the survey. We selected a small probability 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 (e.g., revealing the employee's salary to peers) could be distracting to the employees. Another important feature of the BDM mechanism is that subjects never "lose" money because they choose between receiving money or something else. Many studies use this type of mechanism (Allcott and Kessler, 2019; Fuster et al., 2022), 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).⁷ In order to show that there is some signal among the noise, we present a number of validation exercises. Additionally, we try to mitigate these sources of biases and measurement error 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. And we include a couple of practice questions to familiarize subjects with the BDM elicitation, where we provide feedback if subjects provide values that are abnormally high or abnormally low.⁸ While our efforts may mitigate these biases and measurement errors, they will not eradicate them. This fact is important to keep in mind while interpreting some of the results. For example, while subjects may seem quite heterogeneous in their bids, some of that heterogeneity may simply reflect measurement error in the elicitation exercise. And when we estimate a regression where these variables are on the right hand side of the equation, the measurement error will generate attenuation bias.⁹

3.3. Privacy norms

In the model of Section 2, we assume that privacy norms are stronger for salary information than for seniority information. To provide support for that assumption, we included three subjective questions related to privacy norms. One question, *Unacceptable*, elicits the norm directly, asking whether it is "socially acceptable to ask someone about their salary," with possible answers of

⁷ In some special cases, it may not even be incentive-compatible to report one's true valuation. 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).

⁸ 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 incorrect answers to these questions. However, if the subjects provide responses that are abnormally high or abnormally low, we provide some feedback explaining the typical responses to these questions.

⁹ We also provide a validation of our measures of willingness to pay for privacy and willingness to search for information using the questions on privacy norms described in subSection 3.3, in Appendix A.3.

“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 another question, *Uncomfortable*, to elicit whether the respondent finds it “uncomfortable to ask information about salary 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. A reciprocity norm may augment the effects of privacy norms: individuals may be averse to asking about salary not because they want to avoid bothering others, but because they want to avoid being asked to reciprocate by revealing their own information. To assess this possibility, we ask the question *Reciprocal*: “if you ask a peer about his or her salary, would you expect this peer to ask you about your salary?” The possible answers are “Yes” and “No.”

3.4. Demand for privacy

In the model of Section 2, we assume that the stronger privacy norms for information about salary (relative to seniority) will make it more costly for individuals to acquire information. To provide support for that assumption, we designed a revealed-preference method to elicit preferences for privacy. The survey describes an email that might be sent from the experimenters to five peers, the same five peers whose salaries they have to guess about in the guessing game described below. This email would include the first and last name of the respondent and the respondent's own salary. This email explicitly states that the information is being shared in the context of an experiment. By communicating on behalf of the participant, the channel explicitly shut down the two-way face-to-face communication.

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.¹⁰ 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 choice (either to send or conceal the email). The resulting measure of willingness to pay for privacy can take values from -\$125 to \$125. A positive amount indicates that the subject demands privacy and 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.

The fact that we collaborated with a bank helped to make reward payments smooth and credible. The survey rewards can be deposited directly to the subject's bank account after the survey participation is over.

3.5. Willingness to search for information

We want to elicit the respondents' willingness to search for more information on their own. Our measure of “willingness to search” is meant to capture the individuals' beliefs on the usefulness of having an additional week to gather information. Individuals may be more willing to search because they find it easier to ask peers about their salaries (e.g., due to privacy norms), or because they find it easier to acquire salary information by other means

(e.g., asking the HR manager). However, individuals may be less willing to search simply because they expect that, when asked, their peers' will not provide salary information, or they provide inaccurate information. We measure willingness to search by setting up a guessing game about the salaries of peers. In this instantaneous guessing game, respondents must guess the average salary among a group of peers whose names and position titles are listed within 3 min. 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 additional amount from the experimenter, in addition to the other survey rewards. This extra reward ranges from \$13 to \$65, or approximately between one-half and two days' worth of salary.

We give subjects just three minutes to read the instructions and provide their guess. We chose three minutes so that subjects do not have sufficient time to gather additional information, such as by asking their coworkers. A timer 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. At that stage of the survey, there is no indication that respondents will be able to revise their guess.

Next, we 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 (we also implemented an incentive-compatible elicitation, described below). Next, 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. During the extra week, subjects may consult their peers, or even look up information on the Internet. 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 relates to 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.

We use two alternative measures of willingness to search for information, both of which are motivated by Prediction 2 of the conceptual framework in Section 2. More precisely, Fig. 1.b corresponds to the introduction of additional rewards for accuracy announced at the start of the guessing game, which shifts the MB curve from MB to MB'. If the individual is risk-neutral, the MB curve is shifted upward by exactly the amount of the accuracy reward. Since the individual did not have any time to search for information before providing his or her guess, the individual must stay at q_0^T (in the case of salary information). 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^T , where the MB' curve intersects the MC (Taboo) curve. We use two alternative measures of the willingness to search. The first measure is the increase in the expected probability of winning the game, from q_0^T to q_1^T . Again, if subjects expect the probability of winning the guessing game to be higher with the extra week, that would suggest that they expect to find additional information with the extra time.¹¹

One potential challenge with the first measure of willingness to search is that it is based on probabilities that individuals may not

¹⁰ As before, bids are executed only for a random minority of respondents and remain hypothetical for the rest.

¹¹ Note, we are estimating the marginal cost of increasing accuracy, not the average cost. It is possible, for example, that employees are facing a high marginal cost of searching because that cost curve is convex and they have searched up to a point in which the marginal cost is high. Thus, differences in willingness to search between individuals may reflect differences in marginal cost curves, or differences in their locations on the same marginal cost curve.

report truthfully because they are not incentivized to do so. Thus, we construct a second measure that is also motivated by Fig. 1.b. For this alternative measure of willingness to search, we introduced additional questions designed to be incentive-compatible. After eliciting the probability of winning the guessing game (without and with the extra week), we elicit the subject's willingness to accept to forfeit the right to play the guessing game. Subjects can ask for any dollar value from 0 up to the full reward amount. We explicitly mention that it is in the respondent's best interest to bid their true willingness to accept for giving up the guessing game. As aforementioned, 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, this ratio would reveal the perceived probability of winning the guessing game.

The alternative measure of willingness to search for information is equal to the difference between the normalized willingness to forfeit the guessing game with and without the extra week. In Fig. 1.b, this value is represented by the dark gray area between marginal cost and benefit curves, from q_0^T to q_1^T . 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.¹²

We also designed a validation test for our measure of willingness to search. We randomized the reward for the guessing game to take values $X \in \{\$13, \$26, \$39, \$52, \$65\}$, with equal probability. According to models of rational information acquisition, we would expect employees to be more willing to search for information when the rewards are higher (Woodford, 2001; Mankiw and Reis, 2002; Sims, 2003; Reis, 2006). By randomizing the size of the reward, we generate exogenous variation in the expected benefits of holding accurate beliefs. Thus, we can test if individuals who were randomly allocated to higher reward amounts were indeed more willing to search for information. This validation test can also be illustrated with the framework of Fig. 1. Fig. 1.b shows a hypothetical scenario in which the individual is offered a reward, shifting MB to MB'. When facing a reward, the rational individual responds by searching for more information, up to the point $q_1^T > q_0^T$ (or $q_1 > q_0$ in the case of seniority).

Last, to capture the gross value of the information, we also included in the survey a question that elicits the willingness to pay for information. We implement a variation of the method used in Cullen and Perez-Truglia (2022). We computed a signal about the peer salary (the average salary among a different random sample of five peers), and then let the subjects bid for this piece of information (for more details, see Appendix A.4).

4. Institutional context, data, and subject pool

Our study was conducted in collaboration with a large, private, commercial bank from a lower-middle-income country in Southeast Asia. Due to the sensitive nature of the data described in this paper, the firm has chosen to remain anonymous. However, we can

offer details about the region and scale of the firm: they employ thousands of employees, across hundreds of branches, and manage billions of dollars in assets and revenues.

The firm is typical in many relevant dimensions. First, it does not have an open salary policy. A 2003 survey of Fortune 1,000 companies 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 shows that they are uninformed about salaries and want employers to be more transparent regarding pay (Glassdoor, 2016). Moreover, this firm discourages employees from sharing salary information. Many organizations around the world have similar policies, particularly in the United States (Hegewisch et al., 2011; PayScale, 2018). For example, a 2001 survey of U.S. employees finds that more than one-third work for firms that forbid them from discussing their pay with coworkers (Vault, 2001; Day, 2007).¹³

In our survey, half of the employees report that they do discuss salaries with their coworkers. Similar patterns have been documented in other firms and countries. 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). One natural interpretation for the lack of sharing salary information is that individuals are discouraged by privacy norms. This taboo is believed to exist in a broad range of countries such as the United States (Edwards, 2005), Canada (Bierman and Gely, 2004) and Israel (Fox and Leshem, 2005).¹⁴

The source for the salary taboo is unclear. According to sociologists (Trachtman, 1999), individuals do not feel comfortable discussing pay because it is equivalent to discussing self-worth. Intuitively, salaries reflect things like one's economic contribution to society and whether one is valued by one's employer, which is the type of personal information that individuals often want to conceal from coworkers and even friends and family. Indeed, this conjecture could explain why the salary taboo is so universal: as long as individuals perceive that pay is somewhat related to one's worth, the salary taboo will emerge. Another potential source of the salary taboo is that some individuals (e.g., those with higher salaries) may be hesitant to share their salary information because of the repercussions. For instance, someone may ask them for money, they may be expected to work harder, or their peers may retaliate due to jealousy.

It is also possible that employees do not share information because they are afraid that their employers will punish them if they get caught discussing salaries. While this is a plausible channel, survey data provided by Sun et al. (2021) suggests that it is far from fully explaining the low rates of salary discussions. In a national survey conducted in 2017 and 2018 about half of all workers (48.2%) reported that they were banned or discouraged from sharing their pay, and the other half reported that discussion was permitted or salaries were public. If employer retaliation was the main factor behind the lack of salary discussions, we would expect very high rates of salary discussions in firms where employers cannot punish their employees. On the contrary, Sun et al. (2021) find that even in those firms, close to half of the employees (47.8%) do not share salary information.

¹³ 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).

¹⁴ 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.

¹² 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.

Table 2
Randomization balance test.

	All		Survey Type		Reward Size					P-value
	(1)	Salary (2)	Seniority (3)	P-value (4)	\$13 (5)	\$26 (6)	\$39 (7)	\$52 (8)	\$65 (9)	
Female (=1)	0.73 (0.02)	0.71 (0.02)	0.74 (0.02)	0.50	0.70 (0.04)	0.72 (0.04)	0.73 (0.04)	0.72 (0.03)	0.75 (0.03)	0.89
Age (Years)	29.24 (0.18)	29.39 (0.26)	29.09 (0.25)	0.40	28.80 (0.42)	28.93 (0.38)	29.64 (0.42)	29.23 (0.38)	29.62 (0.41)	0.48
College (=1)	0.86 (0.01)	0.86 (0.02)	0.85 (0.02)	0.76	0.83 (0.03)	0.86 (0.03)	0.85 (0.03)	0.89 (0.02)	0.84 (0.03)	0.55
Seniority (Years)	4.21 (0.13)	4.29 (0.19)	4.13 (0.19)	0.55	3.82 (0.30)	4.02 (0.30)	4.25 (0.33)	4.45 (0.28)	4.44 (0.30)	0.51
Own Salary (Masked)	1.00 (0.02)	0.99 (0.03)	1.01 (0.04)	0.64	0.92 (0.04)	1.00 (0.05)	0.96 (0.05)	1.04 (0.05)	1.06 (0.06)	0.38
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.

On the other hand, employer retaliation does seem to play some role, as 63.2% report sharing salary information in firms where salary discussions are permitted.¹⁵

Another important feature of this context worth discussing is the way in which salaries are set. Several factors can influence whether one employee gets paid more than a peer in this firm. Each position in the company has a pay grade, which is set as part of the bank's overall strategy before a particular individual is assigned to the job. This pay grade sets bounds on the minimum and maximum pay for the position, but these are wide, leaving room for horizontal salary differences. Within a given pay grade, when a new employee is hired, the Human Resources division recommends a salary based on market benchmarks that consider the employee's characteristics, such as his or her experience. However, the employee's manager has the option to override the HR recommendation. For example, the manager can recommend a salary above the market benchmark if the employee has a competing offer. For existing employees, the same rules are at play. HR has policies that influence salary growth, for instance, three years of top performance ratings translate automatically into a salary increase based on a fixed formula. In addition, the manager can also recommend an increase based on the employee's performance during the annual review.¹⁶

4.1. Survey implementation

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.

¹⁵ Relatedly, note that half of the employees who responded to our survey report that they discuss salaries with their coworkers. If employees were significantly concerned about violating the firm's secrecy policy, they would never "admit" this type of behavior in a survey.

¹⁶ Additional details on the institutional context of the firm are provided in Cullen and Perez-Truglia (2022).

We start with the universe of thousands of employees. We focus on two specific units of the firm, with some added filters.¹⁷ We invited the remaining 1,899 employees to take 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.¹⁸ The median respondent took 15 min to complete the survey.

4.2. Descriptive statistics and randomization balance

The subject pool includes employees from 46 different positions, such as tellers, sales associates, and branch directors. Of these, 18% are located in the two headquarter offices, and the rest are scattered across several branches. Table 2 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.

We cross-randomized two features of the survey. In columns (2) and (3) of Table 2, we break down the descriptive statistics by the two survey types, salary and seniority. Column (4) 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 2 provide the correspond-

¹⁷ 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, 2022).

¹⁸ 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).

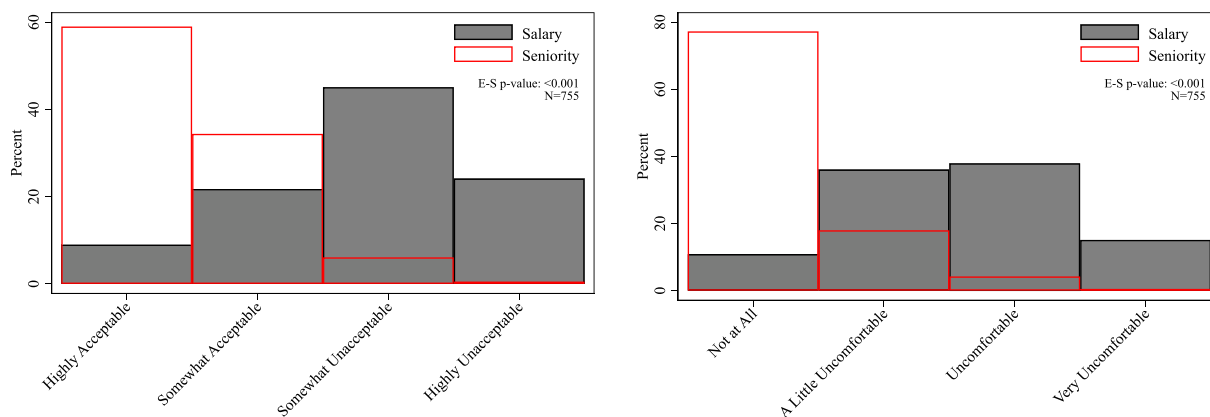


Fig. 2. Privacy norms. *Notes:* Panel (a) shows the distribution of responses to the question *Unacceptable*, asking whether it is “socially acceptable to ask someone about their salary/seniority.” Panel (b) shows the distribution of responses to the question *Uncomfortable*, eliciting whether the respondent finds it “uncomfortable to ask information about salary/seniority to your peers. E-S is the Epps-Singleton characteristic function test of equality of two distributions.”

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

Regarding the pay differentials between peers, 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.

5. Results

5.1. Privacy norms

We start by documenting the degree of privacy norms for salary information and seniority information. Fig. 2 shows the subjects’ responses when asked directly about their privacy norms. This evidence suggests that privacy norms are strong for salary and much weaker for seniority. A whopping 69% of employees find it somewhat or highly unacceptable to ask a coworker about salary, compared to only 7% who find it somewhat or highly unacceptable to ask about seniority. Similarly, 89% of employees find it uncomfortable to ask a coworker’s salary, whereas only 23% find it uncomfortable to ask about seniority. We can also speak to the statistical significance of these differences. In this and other histograms, we report an Epps-Singleton (ES) two-sample test using the empirical characteristic function, which is a version of the Kolmogorov–Smirnov test of equality of distributions that is valid for discrete data (Goerg and Kaiser, 2009). The differences in these distributions across salary and seniority are highly statistically significant (p -value < 0.001). Regarding the economic significance of this difference, we can summarize it using a simple scale. On a scale from 0 to 3, *Unacceptable* averages 1.8 for salary vs. 0.5 for seniority, and *Uncomfortable* averages 1.6 for salary vs. 0.3 for seniority.

The relative unwillingness to ask others about their salaries is consistent with some data from a different, but related, context. In Norway, individuals can use a website to find out the incomes of others. The data on search behavior indicates that the number of searches being conducted was extremely sensitive to the degree of anonymity: When searches for others’ income became non-anonymous, a once active topic of inquiry on the web ceased to be popular (Perez-Truglia, 2020). Last, it is worth noticing that while the norms about privacy are dramatically different, the norms about reciprocity are similar across salary and seniority questions: 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.

5.2. Preference for privacy

Our revealed-preference measure of these preferences for privacy aligns with the direct questions about norms and attitudes. The solid histogram from Fig. 3.a shows the distribution of the willingness to share their 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.¹⁹ Fig. 3.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 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% are not willing to reveal their salaries even under the maximum incentive of \$125. These results together with the ones on privacy norms support our assumption that the cost of acquiring salary information is higher and rises steeply compared to the cost of acquiring seniority information.

While we find significant reluctance to share salaries, our estimates may provide an under-estimate. In our experiment, the firm fully endorsed the sharing of salary information, which made it easier for us to measure how comfortable employees felt discussing salaries when the firm allowed it. In most companies, however, employers discourage this type of behavior, so the unwillingness to share and search for information may be even higher.²⁰ On the other hand, our measured willingness to pay to conceal one’s salary could also constitute an upper bound in some other contexts. For example, individuals could be less reluctant to share

¹⁹ The fact that a minority of people would be willing to pay us to send the informative email 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 to gloat. They may also prefer to share this information when they can avoid reciprocity. Higher-paid employees may be motivated by status concerns to reveal, but nevertheless fear embarrassing lower-paid peers who reciprocate disclosure, for example.

²⁰ Our experiment may also underestimate the reluctance to search for salary information because it may have provided 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.

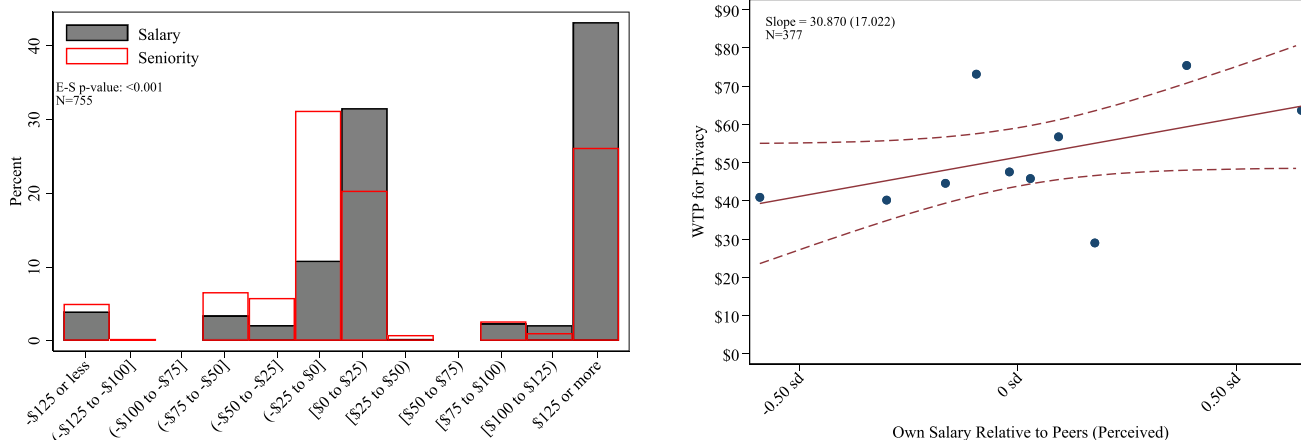


Fig. 3. Willingness to share information with peers. Notes: Panel (a) shows a histogram of the distribution of the willingness to pay for privacy: negative values denote the amount the individual is willing to pay to reveal her information to peers, while positive values denote the compensation the individual is willing to give up in order to conceal this information. Panel (b) provides a binned scatterplot with the relationship between the willingness to pay for privacy and the respondent's perceived relative salary with respect to the reference peer group (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. E-S is the Epps-Singleton characteristic function test of equality of two distributions.

salaries verbally than sending an email as we ask them in our experiment. Also, sharing one's salary without being prompted by a colleague could potentially be more costly than revealing it when being asked, for instance, in a conversation about salaries with other peers.

The hollow histogram from Fig. 3.a is equivalent to the solid histogram but corresponds to the seniority information instead of the salary information. The comparison between those two histograms indicates that individuals are significantly more willing to share their seniority information than their salary information. While 40% of subjects are not willing to reveal their salaries for \$125, the corresponding fraction is about half as large when it comes to seniority information. This significant difference in the distribution of demand for privacy (difference p-value <math><0.001</math>) is consistent with the differences in reported privacy norms.

Note that, even though the privacy norms are much weaker for seniority information, there is still significant willingness to pay to conceal seniority information. That evidence suggests that while privacy norms are a significant barrier to sharing information, it is probably not the only factor at play. For example, the unwillingness to share one's salary information with coworkers may reflect 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, especially if the advantage is not clearly deserved.²¹ If her manager finds out, it may reflect poorly on him or her, and the manager may deny her a raise in the future. 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 salaries 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. Fig. 3.b shows the relationship between the willingness to pay for privacy and the perceived distance between own-salary and the reference group. There is a significant relationship: increasing the individual's perceived relative salary by 1 standard

deviation is associated with an increase in willingness to pay for privacy of \$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 willingness to share seniority information.²² 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, 2022), cause resentment from the manager or even put their relative salary at risk.

This demand for privacy may be an important contributing factor to 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). To complement this evidence on privacy demands, we report some results from a separate survey of 2,033 employees from this same firm (Cullen and Perez-Truglia, 2022). We explained that currently at the firm, salaries are confidential information and asked employees to please report preferences over two alternative disclosure policies. The first policy was described as follows: "the bank created a website showing the average salary by position/unit, for all positions within the bank." We then asked "Would you be in favor or against the creation of a website like this?" Respondents could rate their approval on the following scale: "strongly in favor," "in favor," "I would not care," "against," and "strongly against." The second scenario was described as follows: "The bank created a website with the list of names and salaries of all its employees, including your name and your salary. As a result, you could look up the incomes of any other employee, and any employee could look up your own income." And we asked the respondent whether they would be in favor or against this alternative policy, using the same scale ranging from "Strongly in favor" to "Strongly against." The data from this complementary survey suggests that there is a majority supporting higher salary transparency, but 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.

²¹ In this context salary reflects a combination of performance and good fortune. See Cullen and Perez-Truglia (2019) for evidence of favoritism in this context, as well as a longer discussion about the meritocratic and unmeritocratic sources for pay increases.

²² The results are shown in Appendix A.5.

5.3. Salary misperceptions vs. seniority misperceptions

Relative to the seniority information, employees find the salary information to be more sensitive. In the rest of the analysis, we test if there are, as expected, higher frictions for the more sensitive information (Prediction 1 from Section 2). A natural starting point to measure information frictions is to compare the degree of misperceptions. To measure misperceptions, we compare the employees' guesses about the average salary of peers to the true averages from the administrative records of the firm. The solid histogram from Fig. 4.a shows misperceptions about average peer salary when participants have three minutes to respond. 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 percent error) is 14.6%.²³ These misperceptions are not skewed: approximately as many subjects overestimate the average peer salary as the number of subjects who underestimate it, resulting in an average underestimation of peer salary of just -1.5% (p -value = 0.184).²⁴

Employees understand that their salary perceptions may be inaccurate: on average, they expect a 56% probability to win the guessing game (i.e., that their guess for the average peer salary falls within 5% of the truth). Employees do not seem to be misinformed due to lack of interest. The mean absolute percent error is similar across all the reward conditions, suggesting that subjects were making an active effort to get the right answer even under the lower rewards. And when given the opportunity to buy the signal about the average peer salary, most subjects were willing to pay substantial amounts.²⁵

Fig. 4.b presents the equivalent results to Fig. 4.a but for seniority information instead of salary information. Since there is significantly more variation in peer seniority than in peer salary, we cannot compare the distribution of errors between Figs. 4.a and 4.b. To make the misperceptions more comparable, we can normalize them by the within-group standard deviation of the corresponding outcome (salary or seniority). This measure suggests that individuals are better at guessing peer seniority: the mean absolute error is of 0.707 standard deviations when guessing peer salary but of only 0.469 standard deviations when guessing peer seniority (difference p -value < 0.001).²⁶ These results are in line with the prediction that guesses about salary will be less accurate than guesses about seniority.

Our favorite comparison between salary and seniority misperceptions relies on a benchmark: a "naïve" guess that consists of guessing one's own salary (or seniority). The errors under those hypothetical guesses are shown in the hollow bars from Figs. 4.a and 4.b. Guesses about the average salary of the five peers have a mean absolute percent error of 14.6%. If employees had instead reported their own salary as their guess for the peer average, their mean absolute error would have been 16% instead. The fact that these two misperceptions are similar (difference p -value = 0.126) suggests that employees have little information about salaries besides their own salary.²⁷ In contrast, when it comes to guessing seniority, the mean absolute error for their actual guesses is 29% and it would have been 72% if they had guessed their own seniority

(difference p -value < 0.001). 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.²⁸

5.4. Willingness to search for information

Next we present the results on willingness to search for information. We first report measures of baseline guesses about the salaries of five specific peers, and the respondent's confidence in that guess. Second, we report how employees would change these answers with the option to search, revealing their willingness to search and their expectations of gathering accurate information.

Fig. 5.a shows the distribution of the non-incentivized probability of winning the guessing game with the initial guess. We start by looking at the solid histogram, corresponding to the game without the extra week. Employees understand that their beliefs are not perfect, 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. For that reason, we also present results with our incentive-compatible proxy: the willingness to forfeit the guessing game as a share of the game reward.²⁹ The overconfidence result for the incentive-compatible proxy (shown in Fig. 5.b) is roughly comparable to the results for the non-incentivized measure (shown in Fig. 5.a).³⁰ This overconfidence may be one channel leading employees to under-invest in gathering information. Indeed, participants are similarly overconfident in their ability to predict the seniority of their peers.³¹

What happens when participants consider having an extra week to search for additional information? Employees should expect their accuracy about peer salary to rise with the additional week. We find this hypothesis to be true when we use the non-incentivized or the incentive-compatible measures, shown respectively in panels (a) and (b) of Fig. 5. In both of these panels, the perceived accuracy with the extra week (hollow bars) is shifted to the right relative to the perceived accuracy without the extra week (solid bars). For the non-incentivized measure, the average probability of winning the game about peer salary 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).

²⁸ However, it is worth noting that some employees may be able to "infer" the seniority of some of their peers without necessarily asking them. For instance, if you and your peer started working at the company on a similar date and in the same position, you would immediately infer that peer has your same seniority.

²⁹ 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 non-incentivized measure (correlation coefficient of 0.20, with a p -value < 0.01). We would not expect the correlation to be perfect, however. First, the proxy equals the non-incentivized 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.

³⁰ According to the incentivized measure, on average subjects expect a 43% probability of winning the guessing game, which is still significantly higher than the 32% who actually win it.

³¹ According to the non-incentivized measure, 60% of subjects believe their guess about seniority is within 5% of the truth, and in reality 13% of baseline guesses are within 5% of the truth. Note, there is substantially more variation in peer seniority than peer salary, making it more difficult to guess correctly. For this reason we normalize by standard deviations in the underlying distribution for direct comparisons of misperceptions.

²³ Appendix A.2 provides a finer analysis of the accuracy of the guesses.

²⁴ Despite some methodological differences and using a different subject pool, these results replicate the finding from an earlier study that motivated this work (Cullen and Perez-Truglia, 2022).

²⁵ For more details about the willingness to pay for information, see Appendix A.4.

²⁶ For more details about this comparison, see Appendix A.7. Relatedly, this appendix compares the willingness to pay for information between salary and seniority.

²⁷ Most employees provide a guess different from their own salary. Thus, while they seem to be using information different from their own salary, that information does not seem helpful. For more details, see Appendix A.2.




Fig. 4. Misperceptions: salary and seniority. **Notes:** Histograms of the salary (panel (a), seniority) misperceptions, defined as the difference between the respondent's guess about the average salary/seniority among the sample of five peers (according to the incentivized elicitation) and the actual average salary/seniority (according to the firm's administrative records), divided by the actual average salary. These histograms provide the following benchmarks: what the misperceptions would have been if the respondent had provided a guess equal to her own salary/seniority. E-S is the Epps-Singleton characteristic function test of equality of two distributions.

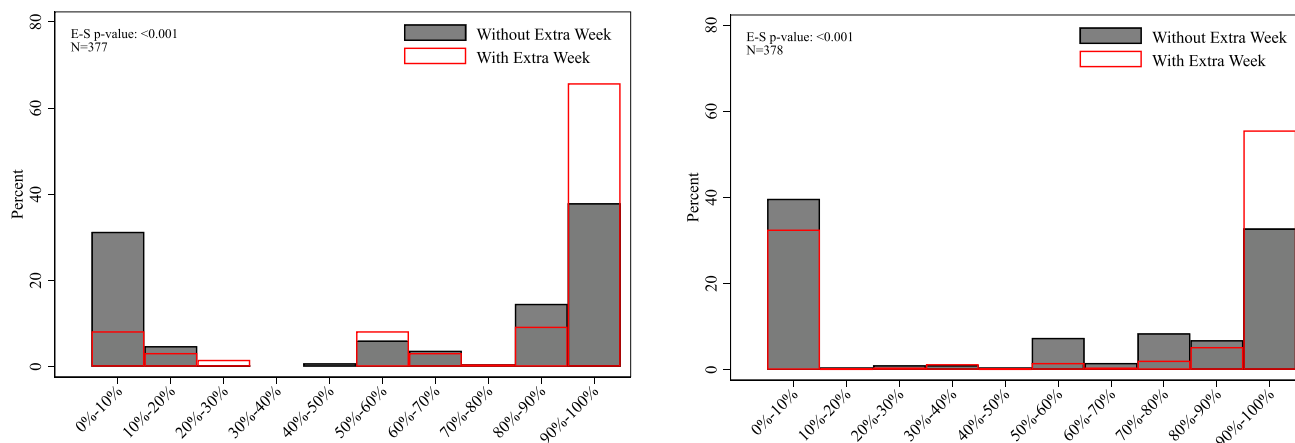


Fig. 5. Salary information: Perceived Accuracy with and without the Extra Week. **Notes:** Panel (a) shows a histogram of the respondent's perceived probability of winning the guessing game, without an extra week to search for information (solid bars) and with the extra week (hollow bars). Panel (b) is equivalent to the first panel, only that instead of using the non-incentivized measure, we use an incentive-compatible proxy: the ratio between the willingness to forfeit the guessing game and the reward amount. E-S is the Epps-Singleton characteristic function test of equality of two distributions.

A standardized measure of search intensity is shown in Fig. 6.a. 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. 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.³² The solid histogram from Fig. 6.a corresponds to the willingness to search for salary information. There is large variation in this measure. The distribution can be roughly divided in three parts. 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 remaining third of the sample is between these two extremes.³³

Next, we present a validation test for our measure of willingness to search. According to Prediction 2 from Section 2), individuals should be more willing to search for information when the gains from doing so are higher. That is, a higher reward in the guessing game should cause individuals to search more intensively. Fig. 6.b presents the results from this test, by comparing the willingness to search with the (randomly assigned) reward amounts. Consistent with our hypothesis, individuals who are assigned to higher rewards expect to search more intensively for salary information. More precisely, a \$100 increase in the guessing reward would result in an expected accuracy increase of 15.5 per-

³² This sample excludes individuals who reported 100% confidence in their initial guesses and for whom there could be no gain in certainty by construction.

³³ 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. In any case, our measure of anticipated gains is the relevant measure for the decision to search for information or not.


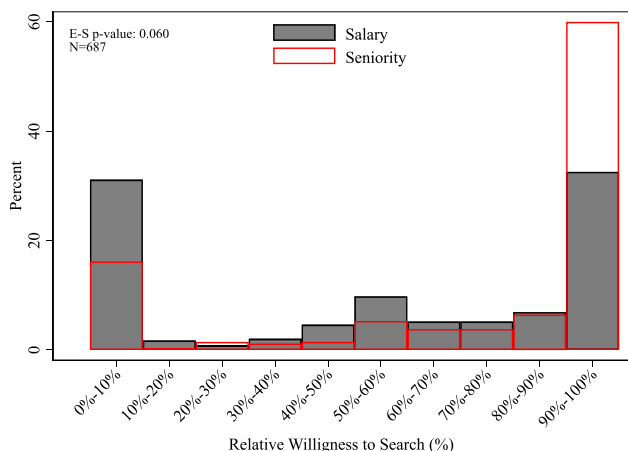


Fig. 6. Willingness to search for information. **Notes:** Panel (a) shows the measure of willingness to search: the difference between the probability of winning the game with and without the extra week, divided by one minus the probability without the extra week (this sample excludes individuals who 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). This relationship is exclusively for individuals being asked about salary information and excludes 30 individuals who were 100% confident in their initial guess. The slope is calculated with ordinary least squares, with robust standard errors reported in parentheses. E-S is the Epps-Singleton characteristic function test of equality of two distributions.

centage points. This difference is not only statistically significant, but also economically large: this 15.5 percentage points increase implies a 28% improvement relative to the average perceived accuracy rate of 56 percentage points.

To gather suggestive evidence on the effects of privacy norms, we focus on the differences in willingness to search for salary and seniority information. In Fig. 6.a, the hollow histogram is similar to the solid histogram only that it corresponds to seniority information instead of salary information. According to Prediction 3 from Section 2, we expect individuals to be more willing to search for information about seniority than about salary. Consistent with this hypothesis, Fig. 6.a shows that employees are not willing to search about salary information nearly as much as they are willing to search for seniority information: on average, the extra week increases the probability of winning the guessing game by 54% for salary and by 75% for seniority. The higher reluctance to search for salary than seniority information are consistent with the difference in privacy norms.

In addition to privacy considerations, there may be other reasons why employees are not willing to search for information, such as the time and effort needed to search.³⁴ Indeed, even for seniority, nearly one-fifth of participants do not expect to improve their accuracy despite the non-sensitive nature of the information. This evidence suggests that the other sources of information frictions, such as attention and time costs, must be significant too.

5.5. Information diffusion

Theoretical and empirical evidence shows that, in presence of information diffusion, individuals who are more central in the network, or more connected to the information source, tend to have significantly lower misperceptions (Alatas et al., 2016; Banerjee et al., 2013). We can use this strategy to test for social learning in the cases of salary and seniority information. As a measure of how connected an employee is to the information source, we measure the overlap between the subject and the five peers he or she has to guess about. This overlap is defined as the time that the subject overlapped in the same position with each of the five peers, as a share of the total time the employee has been working at the

firm. For instance, an overlap of 1 would indicate that the subject has overlapped with each of the five peers during his or her entire tenure at the company. We also construct a measure of how central an employee is in the network. We use data on the emails sent and received by the employees in the months before the experiment. More precisely, we define the employee's centrality as the Eigenvector centrality in the directed graph of emails.

Fig. 7.a shows a binned scatterplot between the subject's degree of misperceptions (measured as the mean absolute error of their guess) and the subject's overlap with the five peers. The red diamonds correspond to the salary misperceptions. The slope between peer overlap and salary misperceptions is negative but close to zero (-0.018), statistically insignificant (p-value = 0.461) and precisely estimated. This slope suggests that increasing the overlap from the minimum possible (0, meaning no overlap with any of the peers) to the maximum possible (1, meaning full overlap with all of the peers) would reduce misperceptions by an statistically insignificant 1.8 percentage points (relative to an average of 15 percentage points). In contrast, the blue circles from Fig. 7.a indicate that the relationship between peer overlap and misperceptions about seniority is negative (-0.133) and highly statistically significant (p-value < 0.001). This slope suggests that increasing the overlap from 0 to 1 would reduce misperceptions by 13.3 percentage points (relative to an average of 29 percentage points).

Fig. 7.b reproduces Fig. 7.a, but using the measure of network centrality instead of peer overlap. The results are robust: salary misperceptions do not change with network centrality, but seniority misperceptions are significantly lower when the subject's centrality is higher. In sum, the evidence from Fig. 7 indicates that seniority misperceptions decline with peer overlap and network centrality but salary misperceptions do not decline with these same variables. This evidence indicates that, consistent with the rest of the evidence presented above, there is significantly more information diffusion for the less sensitive information (seniority) than for the more sensitive information (salary).³⁵

Taken together, the findings presented in the above sections highlight significant differences in the employee's willingness to share and search for information about salary vs. seniority. Our

³⁴ See Sandvik et al. (2020) for an estimate of how much money employees leave on the table by failing to discuss work tactics with their co-workers.

³⁵ In Appendix A.6 we present additional results on the correlation between misperceptions and how heterogeneous salaries are within the peer group, and how different salaries are between the respondent and the peer group.

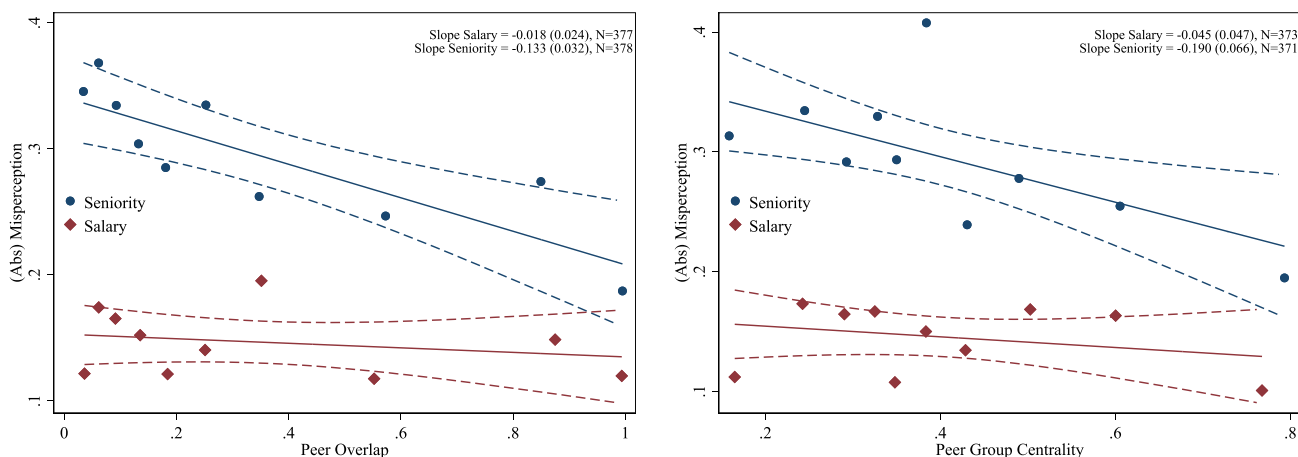


Fig. 7. Salary perceptions and peer connectivity. *Notes:* In panel (a) the binscatter plot and corresponding OLS regression estimates show the relationship between absolute misperception about the average seniority or salary among the five peers (y-axis) and the amount of time that the survey respondent 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. Standard errors are clustered by the unit that the peers and respondent share at the time of the survey. In Panel (b) the specification is identical but the independent variable (x-axis) is the individual's centrality in their peer group measured using the eigenvector centrality of the directed work email graph.

preferred interpretation of these differences in diffusion patterns stems from the difference in privacy norms. However, other channels may contribute to these differences as well. We documented some evidence that strategic reasons to conceal salary information particularly by those who perceive themselves to be the highest paid employees. Another contributing factor may be lingering concerns about the consequences of asking about salary if their manager were to find out. While we think that our communication about this study was effective at alleviating concerns about repercussions, and deliberately did not collect any direct evidence on the act of asking others, it remains a possibility that the general atmosphere of secrecy intimidates employees.

5.6. Gender differences

As a final exercise, we explore whether there are any gender differences in the diffusion of salary information. 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 employees 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 (2021), 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.³⁶

Table 3 presents regressions of several outcomes on a dummy variable that equals 1 if the employee is female and 0 if the employee is male. The evidence from Table 3 suggests that, consistent with the aforementioned survey data from Glassdoor (2016) and Cullen and Pakzad-Hurson (2021), women tend to be less confident than men about their knowledge of peer salaries. In fact,

³⁶ 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.

men and women are equally accurate in their guesses. Column (1) indicates that, according to the non-incentivized measure, the perceived probability of winning the game is 60% for men vs. 54% for women (p-value of the difference = 0.186). Column (2) indicates that, according to the incentive-compatible measure, the perceived probability of winning the game is 50% for men vs. 40% for women (p-value of the difference = 0.041).³⁷ However, the comparison of actual accuracy indicates that these differences in perceived accuracy are misleading. Column (3) of Table 3 indicates that, if anything, women are more accurate than men when it comes to guessing salaries: the share of men winning the guessing game is 26.9% vs. 34.2% for women (p-value of the difference = 0.142); and the mean absolute error of their guesses are also similar by gender (column (4)).

Besides this difference in confidence, there are no significant gender differences in any of the other outcomes reported in Table 3. Female and male respondents feel equally uncomfortable asking their peers about their salaries (column (5)), find it equally unacceptable to ask peers about their salaries (column (6)), and perceive a similar reciprocity norm when asking (column (7)). Employees of both genders are equally willing to reveal their salary to their peers (column (8)), equally willing to search for information (column (9) for non-incentivized measure and column (10) for the incentive-compatible measure). They are also equally likely to pay for readily-available salary information (column (11)).³⁸

6. Conclusions

By designing and conducting a field experiment with 755 employees at a multibillion-dollar corporation, we provide revealed-preference evidence that privacy norms exist and hinder the diffusion of salary information. For example, we show that a significant share of employees are not willing to share their salaries with their coworkers or to ask coworkers about their salaries. When the topic changes from salary information (which has strong privacy norms) to seniority information (which has weaker privacy

³⁷ This difference is no longer significant once we control for the gender composition of the five peers, as shown in Appendix Table A.3.

³⁸ In Appendix A.6 we present these results controlling by gender composition of the peers drawn. We also include heterogeneity results by seniority of the respondent, as a benchmark for the heterogeneity by gender.

Table 3
Salary information: heterogeneity by gender.

	Perceived Acc. (pp)		Accuracy (pp)		Error (pp)		Attitudes		WTP (\$)		Extra Week (pp)		WTP (\$)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
	Direct	Indirect	Actual	Abs. Error	Uncomf.	Unacc.	Recipr.	Privacy	Direct	Indirect	Signal			
Female (=1)	-0.057 (0.043)	-0.107** (0.052)	0.073 (0.050)	0.002 (0.017)	-0.063 (0.099)	0.141 (0.101)	-0.000 (0.033)	-16.765 (12.778)	-0.025 (0.028)	0.017 (0.033)	19.369 (35.612)			
Constant	0.598*** (0.039)	0.504*** (0.043)	0.269*** (0.046)	0.145*** (0.015)	1.472*** (0.083)	1.056*** (0.082)	0.889*** (0.027)	81.599*** (11.794)	0.603*** (0.039)	0.247*** (0.037)	164.821*** (30.623)			

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

norms), the willingness to share and ask for information increases considerably.

In the last decade, several policies have been enacted around the world to promote pay transparency, and new policies are being discussed. Our findings can inform the design of 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 whether employees discuss salaries with their coworkers, at least in the short-term, because employees will continue to be constrained by privacy norms. Indeed, in our own study, despite lifting the non-disclosure policy for the duration of our study, a significant fraction of employees remain unwilling to search for and share salary information. However, if we consider the possibility that social norms are endogenous to the environment created by the firm or government, then policies could have an effect by changing norms over the long run.

Transparency policies around the world differ substantially in whether they disclose sensitive data.³⁹ Our evidence suggests that these policies should avoid compromising employee privacy whenever possible. Our findings also show that, although employees value having access to salary information, they value their privacy too, sometimes even more than the information itself. Moreover, when asked directly, employees strongly support anonymous transparency policies (e.g., disclosing salary averages) and strongly oppose non-anonymous policies (e.g., disclosing individual salaries). For instance, consider the case of U.S. state employees, whose salaries are disclosed on the Internet along with their full names and other personal information. Our findings suggest that this type of non-anonymous disclosure may impose significant costs on some of those employees. The government should consider limiting these costs, such as by anonymizing this information, disclosing average salaries by position and employer, or making individual data available by request only.

Last, our findings suggest an important role for third-party aggregators of salary information, such as Glassdoor, PayScale, and Comparably.⁴⁰ These companies are quickly gaining popularity, which is consistent with our evidence that employees would like to be better informed but have trouble collecting information from coworkers. Our findings suggest that this growing informational sec-

tor may create a lot of value to employees by allowing them to make informed career choices and by improving their leverage in negotiations with employers. Estimating the effects of these salary aggregators is a promising avenue for future research.

Data availability

The data that has been used is confidential.

Declaration of Competing Interest

The collaborating firm provided financial support for the research being conducted. Additionally, Zoë Cullen was a full-time, salaried employee at that firm while the research was being conducted. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.jpubeco.2023.104890>.

References

Acquisti, A., Taylor, C., Wagman, L., 2016. The economics of privacy. *J. Econ. Literat.* 54 (2), 442–492.
 Adjerid, I., Acquisti, A., Brandimarte, L., Loewenstein, G., 2013. Sleights of privacy: framing, disclosures, and the limits of transparency. In: Proceedings of the ninth symposium on usable privacy and security, 1–9.
 Alatas, V., Banerjee, A., Chandrasekhar, A.G., Hanna, R., Olken, B.A., 2016. Network structure and the aggregation of information: theory and evidence from Indonesia. *Am. Econ. Rev.* 106 (7), 1663–1704.
 Allcott, H., Kessler, J.B., 2019. The welfare effects of nudges: a case study of energy use social comparisons. *Am. Econ. J.: Appl. Econ.* 11 (1), 236–276.
 Andersen, S., Harrison, G.W., Lau, M.L., Rutström, E.E., 2006. Elicitation using multiple price list formats. *Exp. Econ.* 9 (4), 383–405.
 Athey, S., Catalini, C., Tucker, C., 2017. The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. NBER Working Paper No. 23488.
 Babcock, L., Laschever, S., 2009. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press.
 Baker, M., Halberstam, Y., Kroft, K., Mas, A., Messacar, D., 2023. Pay transparency and the gender gap. *Am. Econ. J.: Appl. Econ.*, Forthcoming.
 Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O., 2013. The diffusion of microfinance. *Science* 341 (6144).
 Bass, F.M., 1969. A new product growth for model consumer durables. *Manage. Sci.* 15 (5), 215–227.
 Beaman, L., Dillon, A., Lori Beaman, A.D., 2018. Diffusion of agricultural information within social networks: Evidence on gender inequalities from Mali. *J. Dev. Econ.* 133, 147–161.
 Beaman, L., Keleher, N., Magruder, J., 2018. Do job networks disadvantage women? Evidence from a Recruitment Experiment in Malawi. *J. Labor Econ.* 36 (1), 121–157.

³⁹ For example, in 2018, California began requiring that employers provide prospective employees with their current employees' salary ranges (Pender, 2017; Siniscalco et al., 2017). This policy improves employees' bargaining power without compromising employee privacy.

⁴⁰ For more information, refer to <https://www.glassdoor.com/>, <https://www.payscale.com/>, and <https://www.comparably.com/>.

- Bierman, L., Gely, R., 2004. Love, sex and politics? Sure. Salary? No Way: Workplace social norms and the law. *Berkley J. Law Employ.* 25 (1), 167–191.
- Breza, E., Kaur, S., Shamdasani, Y., 2018. The morale effects of pay inequality. *Q. J. Econ.* 133 (2), 611–663.
- Brickley, J.A., Smith, C.W., Zimmerman, J.L., 2007. *Managerial economics and organizational architecture*. McGraw-Hill Irwin, Boston.
- Bursztyjn, L., Ferman, B., Fiorin, S., Kanz, M., Rao, G., 2018. Status goods: experimental evidence from platinum credit cards. *Q. J. Econ.* 133 (3), 1561–1595.
- Bursztyjn, L., Jensen, R., 2015. How does peer pressure affect educational investments? *Q. J. Econ.*
- Card, D., Lemieux, T., Riddell, W.C., 2004. Unions and wage inequality. *J. Labor Res.* 25 (4), 519–559.
- Card, D., Mas, A., Moretti, E., Saez, E., 2012. Inequality at work: the effect of peer salaries on job satisfaction. *Am. Econ. Rev.* 102 (6), 2981–3003.
- Carson, R.T., Groves, T., 2007. Incentive and informational properties of preference questions. *Environ. Resource Econ.* 37 (1), 181–210.
- Cason, T.N., Plott, C.R., 2014. Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *J. Polit. Econ.* 122 (6), 1235–1270.
- Chandrasekhar, A.G., Golub, B., Yang, H., 2019. Signaling, Shame, and Silence in Social Learning. NBER Working Paper No. 25169.
- Charness, G., Gneezy, U., Halladay, B., 2016. Experimental methods: Pay one or pay all. *J. Econ. Behav. Organ.* 131, 141–150.
- Colella, A., Paetzold, R., Zardkoohi, A., Wesson, M., 2007. Exposing pay secrecy. *Acad. Manag. Rev.* 32, 55–71.
- Corbett, W.R., 2002. Waiting for the labor law of the twenty-first century: everything old is new again. *Berkeley J. Employ. Labor Law* 23 (2), 259–306.
- Cullen, Z., Perez-Truglia, R., 2019. The old boys' club: schmoozing and the gender gap. NBER Working Paper No. 26530.
- Cullen, Z., Perez-Truglia, R., 2022. How much does your boss make? The effects of salary comparisons. *J. Polit. Econ.* 130 (3), 766–822.
- Cullen, Z.B., Pakzad-Hurson, B., 2021. Equilibrium effects of pay transparency. NBER Working Paper No. 28903.
- Danziger, L., Katz, E., 1997. Wage secrecy as a social convention. *Econ. Inq.* 35 (1), 59–69.
- Day, N.E., 2007. An investigation into pay communication: Is ignorance bliss? *Personnel Rev.* 36, 739–762.
- Edwards, M.A., 2005. The law and social norms of pay secrecy. *Berkeley J. Employment Labor Law* 26, 41–63.
- Ellison, G., Fudenberg, D., 1995. Word-of-mouth communication and social learning. *Q. J. Econ.* 110 (1), 93–125.
- Fershtman, C., Gneezy, U., Hoffman, M., 2011. Taboos and identity: considering the unthinkable. *Am. Econ. J.: Microecon.* 3 (2), 139–164.
- Fox, S., Leshem, L., 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? *Am. Econ. Rev.* 74 (4), 549–571.
- Fuster, A., Perez-Truglia, R., Wiederholt, M., Zafar, B., 2022. Expectations with endogenous information acquisition: An experimental investigation. *Rev. Econ. Stat.* 104 (5), 1059–1078.
- Gee, K., 2017. Ask me how much money i make. *Wall Street J.*, October 26 2017.
- Gely, R., Bierman, L., 2003. Pay secrecy/confidentiality rules and the National Labor Relations Act. *J. Labor Employ. Law* 6, 120–156.
- Glassdoor, 2016. *Global Salary Transparency Survey*. Technical report.
- Goerg, S.J., Kaiser, J., 2009. Nonparametric testing of distributions—The Epps-Singleton two-sample test using the empirical characteristic function. *Stata J.* 9 (3), 454–465.
- Goldfarb, A., Tucker, C., 2012. Shifts in Privacy Concerns. *Am. Econ. Rev.* 102 (3), 349–353.
- Habyarimana, J., Humphreys, M., Posner, D.N., Weinstein, J.M., 2007. Why does ethnic diversity undermine public goods provision? *Am. Polit. Sci. Rev.* 101 (4), 709–725.
- Hegewisch, A., Williams, C., Drago, R., 2011. *Pay Secrecy and Wage Discrimination*. Institute for Womens Policy Research, June 9 2011.
- Horowitz, J.K., 2006. The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Econ. Lett.* 93 (1), 6–11.
- Karni, E., Safra, Z., 1987. Preference reversal and the observability of preferences by experimental methods. *Econometrica* 55 (3), 675.
- Lawler, E.J., 2003. Pay Practices in Fortune-1,000 Corporations. *World at Work J.* 12 (4), 45–54.
- Mankiw, N.G., Reis, R., 2002. Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips Curve. *Q. J. Econ.* 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? *J. Polit. Econ.* 125 (5), 1683–1721.
- Mobius, M., Phan, T., Szeidl, A., 2015. Treasure hunt: social learning in the field. NBER Working Paper 21014.
- Mobius, M., Rosenblat, T., 2014. Social learning in economics. *Annu. Rev. Econ.* 6 (1), 827–847.
- National Labor Relations Board, 2022. *Your Right to Discuss Wages*. Technical report.
- 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., 2020. The effects of income transparency on well-being: evidence from a natural experiment. *Am. Econ. Rev.* 110, 1019–1054.
- Phillips, M., 2009. Now Comes Lilly Ledbetter. <https://obamawhitehouse.archives.gov/blog/2009/01/25/now-comes-lilly-ledbetter>.
- Reis, R., 2006. Inattentive consumers. *J. Monet. Econ.* 53 (8), 1761–1800.
- Sandvik, J.J., Saouma, R.E., Seeger, N.T., Stanton, C.T., 2020. Workplace knowledge flows. *Quart. J. Econ.* 135 (3), 1635–1680.
- Scott, D., 2003. *Survey of Compensation Policies and Practices*.
- Sims, C., 2003. Implications of rational inattention. *J. Monet. Econ.* 50 (3), 665–690.
- Siniscalco, G.R., Connell, E.M., Smith, C., 2017. State Pay Equity Laws: Where a few go, many may follow. Mimeo.
- Sun, S., Rosenfeld, J., Denice, P., 2021. On the Books, Off the Record: Examining the Effectiveness of Pay Secrecy Laws in the U.S. Institute for Women's Policy Research.
- Trachtman, R., 1999. The money taboo: its effects in everyday life and in the practice of psychotherapy. *Clin. Soc. Work J.* 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.