

Agan, Amanda; Cowgill, Bo; Gee, Laura K.

Working Paper

Salary history and employer demand: Evidence from a two-sided audit

Upjohn Institute Working Paper, No. 22-379

Provided in Cooperation with:

W. E. Upjohn Institute for Employment Research, Kalamazoo, Mich.

Suggested Citation: Agan, Amanda; Cowgill, Bo; Gee, Laura K. (2022) : Salary history and employer demand: Evidence from a two-sided audit, Upjohn Institute Working Paper, No. 22-379, W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, <https://doi.org/10.17848/wp22-379>

This Version is available at:

<https://hdl.handle.net/10419/283974>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

12-2-2022

Salary History and Employer Demand: Evidence from a Two-Sided Audit

Amanda Agan

Rutgers University, amanda.agan@rutgers.edu

Bo Cowgill

Columbia University, bo.cowgill@gsb.columbia.edu

Laura K. Gee

Tufts University, laura.gee@tufts.edu

Upjohn Institute working paper ; 22-379

Citation

Agan, Amanda, Bo Cowgill, and Laura K. Gee. 2024. "Salary History and Employer Demand: Evidence from a Two-Sided Audit." Upjohn Institute Working Paper 22-379. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp22-379>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

Salary History and Employer Demand: Evidence from a Two-Sided Audit

Upjohn Institute Working Paper 22-379

Amanda Agan
Rutgers University
amanda.agan@rutgers.edu

Bo Cowgill
Columbia University
bo.cowgill@gsb.columbia.edu

Laura Gee
Tufts University
laura.gee@tufts.edu

December 2022
Updated February 2024

ABSTRACT

We study how salary disclosures affect employer demand using a field experiment featuring hundreds of recruiters evaluating over 2,000 job applications. We randomize the presence of salary questions and the candidates' disclosures for male and female applicants. Our findings suggest that extra dollars disclosed yield higher salary offers, willingness to pay, and perceptions of outside options by recruiters (all similarly for men and women). Recruiters make negative inferences about the quality and bargaining positions of non-disclosing candidates, though they penalize silent women less.

JEL Classification Codes: J71, M51, C93

Key Words: blinding, disclosure, recruiting, statistical discrimination, salary history, field experiments

Acknowledgments: The authors thank David Abrams, Kyle Coombs, Gordon Dahl, Katie Donovan, Peter Kuhn, Joanna Lahey, Raviv Murciano-Goroff, Bobby Pakzad-Hurson, Patryk Perkowski, Lamar Pierce, David Reiley, Jonah Rockoff, Nina Roussile, Chris Stanton, Olga Stoddard, Olga Shurchov, and Laurina Zhang for valuable feedback and comments, as well as participants at the MEER Conference, the Advances in Field Experiments Conference, the AEA Annual Meeting, the HBS I&I Conference, the Economic Science Association Conference, the Empirical Management Conference, Entrepreneurship and Private Enterprise Development (EPED), NBER Market Design, the NBER Summer Institution (Labor), NBER Summer Institute (Gender), SMS Annual Conference, the Wharton People Analytics Conference, the IZA Workshop, and the seminars at Barnard College, Boston College, Clemson University, Columbia University, Cornell University, the Federal Reserve Board, Florida State University, George Mason University, MIT Sloan, Nuffield College (Oxford), RAND, Rochester University, Strategy Research Forum (SRF), Simon Fraser, SUNY Albany, the University of California Santa Barbara, University of California Santa Cruz, University of Connecticut, University of Missouri, University of Pennsylvania Law School, the University of Texas at Dallas, and Urbana-Champaign (UIUC). We also thank Hailey Brace, Karishma Chouhan, Nadine Fares, Zachary Finn, Matt Fondy, and Norman Yuan for research assistance. We would also like to thank the anonymous referees and the editor for their thoughtful and helpful comments. We acknowledge financial support from the W.E. Upjohn Institute, Facebook Research and the Kauffman Foundation. The experiments in this paper were approved by an Institutional Review Board and preregistered at the AEA RCT Registry under AEARCTR-0003088.

1 Introduction

Labor markets are rife with asymmetric information. Knowing candidates’ salary history can reduce uncertainty and shape workforce composition and income. However, the meaning of applicant salaries can vary. In several important economic models, workers are primarily differentiated by their outside options (which are unobservable to employers), such as in models of monopsony ([Manning, 2003](#); [Card et al., 2018](#); [Jäger et al., forthcoming](#); [Caldwell and Danieli, Forthcoming](#)), efficiency wages ([Lazear et al., 2016](#)) and some models of on-the-job search ([Burdett, 1978](#)). In these models, salary history could signal workers’ outside options and bargaining positions. By contrast, an additional set of influential models features workers differentiated by “hidden ability” ([Gibbons and Murphy, 1992](#); [Oyer and Schaefer, 2011](#); [Waldman, 2013](#); [DeVaro and Kauhanen, 2016](#)). Talent is often hard to measure, and salaries could reveal latent aptitude.

What information does a worker’s salary reveal in practice? What message is sent by withholding this information? Does it matter if the disclosure is prompted by a salary history question or provided voluntarily? These questions are particularly relevant for female and/or minority job-seekers—for victims of discrimination, disclosing past wages could anchor today’s salary negotiations around historical patterns.

In this paper, we develop a conceptual framework around the effects of salary disclosure. Then, we execute a novel field experiment in a real world setting to understand how salary disclosure impacts employer demand for workers. In our experiment, we assumed the role of an employer and hired hundreds of recruiters to make decisions on over 2,000 job applications for a software engineering position. The recruiters in our field experiment were paid real wages and faced real incentives but were not informed that our job applicants and job openings were fictitious.

Our “two-sided” audit study varies attributes of both sides of the market, the candidate and the employer. For candidates, we randomly vary whether candidates disclose their current salaries or withhold them as well as the gender of the candidate; on the employer side, we randomly vary the presence of a salary history prompt on the job application form. Among candidates who disclose, we also vary the levels of salaries disclosed, with an average male-favoring gender wage gap. Recruiters are then asked to evaluate our candidates and provide information such as the suggested salary offer, the maximum salary the firm should pay the candidate (e.g. willingness-to-pay, or WTP), and whether the candidate should be called back. Our design allows us to observe recruiter’s beliefs about all job candidates, whether or not they were recommended for a callback by the recruiter.

We find that when a worker discloses a higher salary (rather than a lower one), recruiters raise their beliefs about the worker’s latent quality and outside options, and hence the salary offer they propose for the worker. Every \$1.00 increase in disclosed salary increases i) beliefs about the median outside option by \$0.73, ii) employer WTP by \$0.64, and iii) the salary offer by \$0.66.

Across several outcomes, we find no statistically significant differences in recruiters’ assessments in prompted (versus unprompted) salary history disclosures. The lack of prompt effects in our results does not mean that salary history prompts do not matter. Rather, it suggests that the prompts matter via their impact on a candidate’s choice to disclose (or not), rather than through the employer’s interpretation of disclosures (when they happen). In a companion paper surveying job candidates, we find that prompting increases a candidate’s willingness to disclose, particularly for women (Cowgill et al., 2022).¹

Do the recruiters interpret salary disclosures differently by the candidate’s gender? Conditional on disclosing, we find no gender differences in the effects of additional dollars of disclosed salary on recruiter assessments. However, we do find gender differences in the fixed effect of choosing to disclose at all (rather than staying silent). Men who disclose an average salary receive salary offers \$9,400 higher than men who do not disclose. For women it is only \$2,300 higher. These differences raise questions about the meaning of silence. On average, when candidates decline to disclose their salary history recruiters’ lower beliefs about both the value of the candidate (WTP) and the candidate’s outside options. Unsurprisingly, these silent candidates are also given lower salary offers.

Our results suggest that recruiters anticipate positive selection into disclosing. However, we also see that recruiters do not punish silent candidates as much as they could. In fact, some workers who disclose low amounts are regarded as even worse than non-disclosers (they might have been better off silent). Why are recruiters not more punitive? Recruiters may believe that disclosure choices are affected by factors unrelated to a worker’s value. In fact, we find that silence is less damaging for female candidates compared to males. Our companion survey suggests that women are less willing to disclose their salaries, even after controlling for current salary and other characteristics (Cowgill et al., 2022). Recruiters may anticipate this, and realize that female non-disclosure may be less informative. The flip side of this recruiter behavior is that the benefit of disclosing is also smaller for women.

To fully reap the benefits of potentially higher salary offers, a candidate must be selected for a callback (and an offer). How do salary disclosures affect recruiters’ callback recommendations? Our conceptual framework shows that disclosures could impact a recruiter’s

¹Cowgill et al. (2022) is an extension of Agan et al. (2020) and contains a second round of survey data collection and an expanded set of questions and results.

decision to call a candidate back in two contrasting ways. On the one hand, if disclosures mainly reveal outside options (and little about candidate quality), higher salaries could deter callbacks, like a buyer reducing their purchase probability due to a higher price. This aligns with economic models where higher prices decrease surplus for employers. On the other hand, if disclosures signal strong candidate quality more than outside options, they could boost callbacks. An exceptional candidate could cost more, but could bring in even more value. Our data suggest that salary disclosures tend to influence perceptions of outside options more than perceptions of quality, suggesting a potential decrease in callbacks with higher salaries. However, our empirical evidence for whether higher salaries actually decrease callback probabilities for candidates is noisy (in part because of the binary nature of the callback variable).

This paper provides several contributions to the literature. We contribute a new application of disclosure theory and statistical discrimination, and new field experimental evidence about how voluntary disclosures are understood by employers. Our paper adds to the literature on “price as a signal of quality” in a labor market setting (akin to [Roussille, forthcoming](#)). We also develop the idea of “price as a signal of outside options.” In thin and/or monopsonistic markets (e.g. [Manning, 2003](#); [Ashenfelter et al., 2010](#); [Caldwell and Danieli, Forthcoming](#)), beliefs about outside options can affect wages through a separate, non-quality channel. We formalize this notion in our model, and our empirical results suggest that beliefs about outside options are more affected by disclosures.

A classic literature examines learning and strategy around voluntarily disclosed information ([Viscusi, 1978](#); [Grossman and Hart, 1980](#); [Milgrom, 1981](#); [Grossman, 1981](#)). In most disclosure models, agents have a unidimensional hidden characteristic, and the audience’s preferences are monotonic in this characteristic (“more is better” or “less is better.”) In these models, voluntary silence (refusal) is viewed as a negative signal of quality, and this leads to full revelation. Empirically, unraveling is not always observed ([Mathios, 2000](#); [Dranove and Jin, 2010](#); [Jin et al., 2021](#)). Hiring settings may have distinct considerations. Candidates could have multiple dimensions of latent qualities: “hidden talent” and “hidden outside options.” Disclosures could reveal information about both, either directly or through strategic considerations. In addition, a salary disclosure could be “too high” as well as “too low,” rather than more is always better (or less). Our conceptual framework integrates these considerations together theoretically, while our field experiment studies them empirically.

We also provide new results about gender differences in job search and how employers anticipate and react to them. We propose microfoundations for these differences in our setting, and we trace how these foundations affect candidate choices and employer reactions. In our empirical results, we find that recruiters incorporate expectations about gender differences

into their decisions.

We develop a new experimental methodology to support the research questions around this topic. Methodologically, our work is related to recent innovations in correspondence audit methodology ([Bartos et al., 2016](#); [Kessler et al., 2019](#); [Avivi et al., 2021](#)). Our two-sided design allows researchers to collect detailed outcome data beyond the binary callback choice that reveals and suggests underlying mechanisms and to experimentally manipulate employer screening policies. Unlike traditional audit studies, in which subjects are not compensated for evaluating fictitious candidates, our recruiters are compensated at their normal pay rate in a natural way through standard hiring practices.

Prior research suggests women are less aggressive in job search behavior, have lower propensity to enter competitive environments ([Niederle and Vesterlund, 2007](#); [Flory et al., 2015](#)), self-promote less ([Exley and Kessler, 2022](#); [Murciano-Goroff, 2021](#)), ask for lower salaries from employers ([Roussille, forthcoming](#)), and are less willing to disclose salaries ([Goldfarb and Tucker, 2012](#); [Cowgill et al., 2022](#)).² We propose a microfoundation: there are differences in the psychological costs of disclosure that are correlated with gender, but not correlated with other characteristics such as talent.³ We then draw out the implications for the demand side: treating the same negotiation signals differently by gender is necessary to update beliefs accurately. Our model connects gendered negotiation behavior to theories of voluntary disclosure, employer learning, and unraveling.

Our work relates to policies that obscure certain personal details of job applicants from decision-makers. Recent laws limit credit checks, drug tests, gender questions, criminal history questions and other personal details.⁴ In many settings, concealing information requires cooperation from the supply side, which can circumvent concealment by disclosing voluntarily. Through unraveling dynamics, private information could become fully revealed ([Viscusi, 1978](#); [Grossman and Hart, 1980](#); [Milgrom, 1981](#); [Grossman, 1981](#); [Jin et al., 2021](#)). Although the ideas in our paper could apply to any ban on information-seeking, we specifically

²Other examples include gender differences in the propensity to apply for a job given the number of other applicants ([Gee, 2019](#)), the choice to disclose skills ([Murciano-Goroff, 2021](#)), the perceived returns to job search ([Adams-Prassl et al., 2023](#)), and the choice to negotiate wages ([Laschever and Babcock, 2003](#); [Biasi and Sarsons, 2022](#)).

³Although disclosure costs can encompass many things, they are distinct from other theoretical explanations for differences in negotiation behavior — for example, the theories that men enjoy competition more ([Niederle and Vesterlund, 2007](#)), that one gender has more biased beliefs about its own abilities ([Bordalo et al., 2019](#)) or that genders vary by risk aversion ([Croson and Gneezy, 2009](#); [Marianne, 2011](#)). In our framing, the act of disclosing enters workers' utility functions directly. [Exley et al. \(2020\)](#) similarly studies negotiation costs by gender, including indirect costs of unsuccessful negotiations.

⁴[Bartik and Nelson \(2016\)](#); [Friedberg et al. \(2021\)](#); [Corbae and Glover \(forthcoming\)](#) study credit checks, [Card et al. \(2021\)](#); [Kuhn and Shen \(2023\)](#) study gender information, and [Doleac and Hansen \(2020\)](#); [Agan and Starr \(2018\)](#) study criminal history questions.

study salary history. A nascent literature directly studies policies that ban salary history information (Bessen et al., 2020; Davis et al., 2020; Hansen and McNichols, 2020; Mask, 2023; Sinha, 2020; Sran et al., 2020). A few other researchers have examined the effect of salary disclosures and salary history bans using experiments in online markets (Barach and Horton, 2021), laboratory settings (Khanna, 2020), or in real-life educational institutions (Sherman et al., 2023). Our field experiment uses recruiters for corporate jobs, and is focused on the mechanisms underlying these policies (voluntary disclosure, unraveling, and prices signaling both quality and outside options). Our results address design considerations around policies hoping to insulate decision-making from personal details.

In that sense, our paper provides suggestive evidence about the potential impacts of salary history bans. Survey evidence suggests that up to 43% of job applicants are asked about salary history during a job search (Hall and Krueger, 2012; Barach and Horton, 2021; Cowgill et al., 2022). Our own survey evidence shows that the most common method of inquiring about salary history was in writing (on job application forms) — among workers who were asked, 45% were asked this way.⁵ Written salary history questions on job application forms are so common, in fact, that some jurisdictions explicitly address the practice in the text of legislation (HRDive.com). Job interviews (34%) were the second most common context. As of January 2023, 21 states and 21 local jurisdictions have adopted some form of salary history bans (HRDive.com). In 2019 and 2021, a federal salary history ban passed the House of Representatives. These laws vary in their details, but nearly always prohibit oral or written questions about salary history, even if the questions are posed as optional. However, applicants under the bans are still permitted to *voluntarily* disclose salary history information *without prompting*. Our own survey finds that 52% of job seekers only disclose if asked, while 28% always disclose and the remaining 19% never disclose (Cowgill et al., 2022). In most jurisdictions, employers are allowed to use or confirm voluntarily disclosed information.

The rest of the paper is organized as follows: Section 2 describes a conceptual framework of employer updating from salary information. Section 3 lays out our experimental design and Section 4 outlines our empirical specifications. Our experimental results are in Section 5. Section 6 discusses the generalizability of our findings and implications for bans, and the final section concludes.

⁵We ran a smaller survey of about 50 software engineers and found that one-third of software engineers said they had been asked for salary history, and of those one-fifth stated they were asked on the application form.

2 Conceptual Framework

In this section we present a simple theoretical framework. This model is composed of two blocks: a model of employer learning based on voluntary disclosure, and a model of firm-specific hiring and wages in an imperfectly competitive labor market (Manning, 2003; Card, 2022). Both blocks adopt typical conventions of prior models, and our joint framework straightforwardly merges the blocks. Our aim is to use the combined model to show how an employer’s hiring and wage choices are influenced by voluntary disclosures in an imperfectly competitive market.

Utilities. The two players are an employer and a job applicant who has applied to the employer. The sequence of the game is simple: The candidate discloses a privately-known characteristic (such as salary history) or remains silent, and then the employer makes callback and wage choices. The employer’s choice of callback is $b \in \{0, 1\}$ and the salary offer (for those called back) is $s \in \mathbb{R}$. Making a callback costs the employer c . Candidates who accept the employer’s offer generate a payoff v for the employer. The employer’s utility is thus $(v - s - c)$ if the worker is hired, $-c$ if the employer makes an offer that is rejected, and zero otherwise.

The applicant gains some utility α from getting a callback from the employer (commonly known). In addition, the candidate gets more utility from offers that are better than their outside option. Let η equal the salary necessary for the employer to outbid the candidate’s next-best option. The candidate’s next best option could include working at a different firm, leisure, or providing childcare or other unpaid household work; η is simply the amount necessary to outbid the utility of this outside option. η is privately known to the worker, and unknown to the employer though they may form a prior about it. Finally, the worker may have disclosure costs of m if he makes a non-empty report in the disclosure part of the model. We use salary history as an example, but in principle the report could be about other privately-known personal characteristics. Together, the candidate’s utility is thus:

$$u(b, s; \eta) = \underbrace{\eta}_{\text{Outside option}} + \underbrace{\alpha \cdot b}_{\text{Additional payoff from a callback}} + \underbrace{\beta \cdot \mathbb{1}(s > \eta)(s - \eta)b}_{\text{Additional payoff from callback w/ salary above outside option}} - \underbrace{m \mathbb{1}(r \neq \emptyset)}_{\text{Disclosure costs}} \quad (1)$$

Information. The employer has a joint distribution of beliefs about the candidate’s value, outside options and salary history $F(v, \eta, h)$, and these beliefs can be updated after the candidate’s revelations in the disclosure part of the model. α and β are known to the employer.

Candidates also have a privately-known variable $h \in [\underline{h}, \bar{h}]$ (e.g., salary history) that can either be disclosed or not. h can be verified if disclosed, and thus the candidate cannot lie (in our example, employers could verify salary history with pay-stubs or bank statements). The candidate's report space is $r \in \{h, \emptyset\}$. When undertaken, disclosure costs the candidate a privately known amount m , independently drawn from a publicly-known distribution G_m .

To help understand the effects of disclosure, we make two additional assumptions. Let $v(h) = E_F[v|h]$, or the candidate's expected v with a history h , and let $\eta(h) = E_F[\eta|h]$ equal the expected outside option η for a candidate with history h .

Assumption 1 (Informativeness). *A candidate's expected value $v(h)$ and outside options $\eta(h)$ are both weakly increasing in the signal h (e.g., $\frac{\partial v(h)}{\partial h} \geq 0$ and $\frac{\partial \eta(h)}{\partial h} \geq 0$).*

This assumption means salary history is a weakly positive — but potentially noisy — signal of value v and outside options η . For the next assumption, let $\pi(h) = E[v - \eta|h]$, or the expected employer surplus for a candidate with history h .

Assumption 2 (Monotonicity). *Expected employer surplus $\pi(h)$ — is weakly monotonic in the signal h (e.g., $\frac{\partial \pi(h)}{\partial h} \geq 0$, or $\frac{\partial \pi(h)}{\partial h} \leq 0, \forall h$).*

The direction of the monotonicity could in theory go either way. On average, employer surplus either rises with salary history — or it could fall with salary history. We will later discuss either scenario with examples, and show how the direction of these relationships affects our theoretical predictions. Our experiment measures this relationship empirically in a natural setting.

2.1 Hiring and Wages

We approach the game backwards starting with the hiring and wage-setting block. Given a joint distribution of posterior beliefs $F(v, \eta, r)$ about the candidate's value and outside options (given a disclosure action $r \in \{h, \emptyset\}$), the employer can calculate a TIOLI (take-it-or-leave-it) offer s^* :

$$s^* = \underset{s}{\operatorname{argmax}} \iint \underbrace{\mathbb{1}(s > \eta)}_{\text{Whether candidate accepts}} \cdot \underbrace{(v - s)}_{\text{Net value of employment, if accepts}} \cdot \underbrace{f(v, \eta, r)}_{\text{Joint probability}} dv d\eta \quad (2)$$

Given this s^* , the employer can then decide whether to extend a callback at all. The employer will extend a callback if:

$$b^* = \iint \underbrace{\mathbb{1}(s^* > \eta)}_{\text{Whether candidate accepts the optimal TIOLI offer } s^*} \cdot \underbrace{(v - s^*)}_{\text{Net value of employment, if accepts the optimal TIOLI offer } s^*} \cdot \underbrace{f(v, \eta, r)}_{\text{Joint probability}} dv d\eta > \underbrace{c}_{\text{Fixed cost of a callback}} \quad (3)$$

Higher beliefs about the candidate's value v increase the employer's returns from giving a callback and for a generous TIOLI amount s^* . However, higher beliefs about outside options η increase s^* , but decrease returns of sending a callback due to lower net values of acceptance (unless beliefs about v simultaneously increase).⁶

From the candidate's perspective, these equations link the choice of callback to the salary necessary to recruit the candidate. The candidate wants a higher salary, but a salary that is "too high" may deter the employer from a callback. This is true for all candidates, but particularly for $\alpha > 0$ candidates who enjoy utility from low offers.

2.2 Disclosure

From here, we have the ingredients to study the candidate's disclosure choice in the first part of the model. Because candidates cannot lie, employer learning is straightforward when a candidate discloses. However, inferences about silent candidates depend on the parameters of the game. The theoretical predictions from our framework are therefore ambiguous, but our framework helps show what the predictions depend on (and focus empirical tests towards these parameters). Below we mention three cases.

Case 1: Increasing employer surplus $\left(\frac{\partial \pi(h)}{\partial h} \geq 0 \right)$. If higher salary history h candidates provide higher surplus ("bargains at the top"), then there is less tension between a high salary history and becoming too expensive. Unraveling proceeds typically: The candidate with the highest salary discloses, as long as disclosure costs m are not too high. The candidate with the second highest salary decides similarly, rather than be pooled with the others, and so on. The only obstacle to full unraveling is the disclosure costs m . Because of these, silent workers contain a mixture of candidates with low salaries (strategically withheld), and candidates with high costs.

⁶By higher beliefs about v , we mean a new, joint distribution of beliefs $F(v, \eta, r)$ in which the marginal distribution of v first order stochastically dominates the original set of beliefs (and the same for η).

Case 2: Decreasing employer surplus $\left(\frac{\partial \pi(h)}{\partial h} \leq 0\right)$, **large β** . Salary history could also be *negatively* correlated with employer surplus (“bargains at the bottom”). In this case, employers will prefer to target low salaries. However, unraveling does not necessarily proceed in reverse. This depends on how much workers benefit from extra salary (β) versus the flat benefits from an offer (α). As β becomes larger compared to α , the candidate cares only about high salary when he is called back. The strategy is to disclose if high, hide if low. The standard unraveling logic proceeds. Silent workers again contain a mixture of low salaries and high disclosure costs.

Case 3: Decreasing employer surplus $\left(\frac{\partial \pi(h)}{\partial h} \leq 0\right)$, **large α** . If α becomes larger relative to β , the candidate cares only about getting an offer. The best strategy then is to disclose if low, hide if high. The unraveling logic now proceeds *in reverse*: low salaries are disclosed, and high salaries are hidden. Silent workers contain a mixture of candidates with high salaries or high disclosure costs. For intermediate values of α and β , both high and low can be bad. As a result, silence contains high cost candidates, and a mixture of high and low h candidates. The exact mixture depends on the level of α and β .

Extension: Observable Characteristics. The distributions in our framework could differ by observable characteristics such as gender. Employers could thus treat candidates who disclose the same report r differently by gender. Different disclosure costs m may be relevant to gender differences. Across multiple studies, women are less comfortable disclosing their salaries ([Goldfarb and Tucker, 2012](#); [Cowgill et al., 2022](#)). High disclosure costs for women suggest that their silence contains less information.

2.3 Key Parameters that Shape Results

Our model shows why the effects of disclosure and silence are theoretically ambiguous. Rather than having mechanical effects on hiring and wages in any direction, the interpretation of these signals depends on the economic environment.

The value of a low offer (α). A worker who only cares about his salary at this employer may have $\alpha = 0$. However, some job-seekers may have value for being called back, even at a low salary. An outside option may still be a useful bargaining chip with current employers, even if it’s for a slightly lower salary. Second, a worker may value non-pecuniary aspects of the job ([Stern, 2004](#)). Idiosyncratic tastes for employers – separately of wages – features in

many existing models ([Manning, 2003](#)). Finally, a low starting salary could be compensated for with greater wage growth in the future, either at the same company or by switching. Through any of these mechanisms, a callback might be useful even if the salary offer is not high.

How informative is h ? $\frac{\partial v(h)}{\partial h}$ and $\frac{\partial \eta(h)}{\partial h}$. Our framework says that salary history is a noisy signal of quality and outside options. However, the degree of noise in the signal could range anywhere from quite low to very high. As salary history signals become more noisy, all effects attenuate.

Bargains at the Top? $\frac{\partial \pi(h)}{\partial h} \geq 0$. In some labor markets, candidates with high salary histories are more expensive, but are “worth every penny” because they deliver even higher value. “Bargains at the top” appear in many markets for goods and services. A common reason is that few buyers can afford the most expensive options. As a result, the market for expensive items is relatively thin; thinness prevents sellers from commanding 100% of their full value. “Bargains at the top” are also possible if buyers (employers) evaluate expensive options idiosyncratically. Finally, price controls could also generate surplus at the top. The most expensive rent-controlled apartment in NYC may be a bargain, because the price ceiling prevents the seller from extracting the buyer’s full WTP.⁷ Elite athletes are sometimes said to be “underpaid” because salary caps prevent stars from seeking their full value ([NPR, 2018](#)). Scenarios like these resemble Case 1: High h candidates have nothing to lose from disclosing and pressure lower h to disclose.

Bargains at the Bottom? $\frac{\partial \pi(h)}{\partial h} \leq 0$. The relationship could also go in the opposite direction: The greatest bargains may appear where salary history h is low. Even if high h ’s are higher quality, they may also command higher salaries (eating into employer surplus). Similarly to above, employers may value low h ’s idiosyncratically, leading to little competition at the bottom. The labor market for entry level positions is thin ([Pallais, 2014](#)), in part because of uncertainty about their quality. “Bargains at the bottom” create scenarios like Cases 2 and 3. If workers care a lot about getting an offer – and are possibly willing to trade-off salary – then low h will be most tempted to disclose (Case 3). If they care only about salary, then unraveling proceeds as before (Case 2).

⁷In some settings, price floors could similarly diminish deals on the lower end of the market, by increasing prices at the low end to levels approaching values ([Horton, Conditionally Accepted](#)). This would potentially leave more buyer surplus on the higher end of the market.

3 Experimental Design

Our framework implies that the effect of salary disclosures on hiring outcomes is ambiguous and depends on the specific characteristics of the setting. To examine the empirical impact of salary disclosure, we implemented a two-sided audit study. We hired and paid a recruiting workforce of 256 (real) recruiters to screen 2,048 (fictitious) job applications on behalf of a (fictitious) firm. This design allows us to vary both characteristics of the candidates and characteristics of the firm. We focus on the software engineering industry, in part because it helps place our results in the context of wider labor market trends (Appendix A). In a later section (6), we discuss the generalizability of results from this setting.

3.1 Our Recruiting Workforce

Our design uses the practice of delegating recruitment decisions to specialists. In the past two decades, the delegation of recruiting to third parties has become widespread (Landay and DeArmond, 2018; Black et al., 2020; Cappelli, 2019; Cowgill and Perkowski, forthcoming). Firms that use this practice either hire individual recruiters on a temporary, contract basis, or they outsource recruiting entirely to a third party organization.

To staff our recruiting workforce, we identified recruiters who are typical of those hired by companies through the recruiting industry and engaged them in a natural way. The recruiters in our experiment appeared on LinkedIn and UpWork offering recruiting services (both freelance and full-time), and we directed them through UpWork for the experiment’s payroll needs. We identified and contacted professional recruiters as outlined in Appendix B. We only contacted recruiters who had prior recruiting experience and a U.S.-based location. We offered to pay recruiters their hourly rate as it appeared on their profile. The human resource workers in our experiment were similar to those in the U.S. as a whole (Appendix C).

3.2 Our Recruiting Task

Recruiters were given three sets of documents: 1) information about the software engineering job they were hiring for, 2) eight one-page job candidate applications, and 3) a structured evaluation form to provide feedback about the candidates. These structured evaluations are typical in the industry (Black et al., 2020; Jobvite, 2015; Monster, 2018) and we detail the items in our version below; a full copy of the structured evaluation is available in Appendix L.

Before they were hired, recruiters were given a general sense of the information they would be asked to provide via the task description, but were not given the specific questions in advance. Each question in the structured assessment appeared on a separate page. Recruiters were required to answer each question about all eight candidates before proceeding to the next question. No recruiters sought to revise their earlier answers. Recruiters also saw text fields on each page for additional comments on the items above and below. We monitored these comments (and all other communication) for any messages that would change our interpretation of the evaluations. Note that we describe recruiter choices using economic terms, but our feedback form (and all communications with subjects) used everyday business language.

Callback. Like a traditional audit study, for each candidate we ask the recruiter whether they recommend that candidate for an interview. Recruiters were allowed to suggest as many candidates as they deemed a good fit for an interview, which is common in labor settings featuring high demand for qualified workers. We conceptualize this choice as the callback (b^*) modeled in Equation 3 of our framework.

Salary Offers. Recruiters were also asked to suggest a take-it-or-leave-it (TIOLI) salary offer for each job candidate. Hall and Krueger (2012) find that two-thirds of workers report believing that the offer they were made by an employer was a take-it-or-leave-it offer. We conceptualize these as the s^* choice modeled in Equation 2 of our framework. The instructions stated that the employer currently paid between \$70,000 and \$120,000 in salary annually for this role; however, the recruiter was allowed to suggest differently. We observe a TIOLI offer even if the recruiter did not believe the candidate should be called back. To observe salary offers even for those not suggested for a callback, recruiters were told, “For candidates you do NOT suggest interviewing, please enter the amount you think they should be offered were they to pass an interview — this may be helpful for us in the future.”⁸ The offer and callback decisions were on the first page of our evaluation.

Willingness to Pay (WTP). Recruiters are sometimes asked to evaluate an offer or a proposal by the candidate. In these cases, recruiters need to decide (in everyday language) “how high they’re willing to go.” According to Barach and Horton’s 2021 survey, employers make the first offer about 60% of the time. In the remaining 40% of cases, a job candidate makes

⁸Recruiters made recommendations about the annual base salary of compensation only, although the firm instructions said, “We also offer benefits including health insurance, stock, and performance-based annual bonuses,” without specifying their amounts.

a first offer. As a result, our recruiters were also asked to report the maximum salary offer from the candidate that the firm should accept.⁹ We conceptualize these as the v (expected value) that appears in our framework. By reporting a threshold, we observe the recruiter's value for the candidate. Reporting any threshold below their true value is a dominated strategy (Becker et al., 1964).

Outside Option Distributions. For each candidate, recruiters were asked to state the TIOLI salary offers that they believe the candidate would accept. We specifically asked for salary offers the candidate would accept with 95% probability, 50% probability and 5% probability. We interpret these as the salary necessary to outbid the candidate's outside option (η in our conceptual framework), at the 5th, 50th and 95th percentiles of the recruiter's belief distribution. The candidate's outside option could include working at a different firm, leisure, or providing childcare or other unpaid household work. In the main text we only report results on the 50th percentile, and for brevity we call this "outside option" in our tables. In the Appendices we also report results for the 5th and 95th percentiles.

In addition to the assessments above, we asked a few additional questions as explained in Appendix section E.

The goals for a recruiting task such as ours are to identify candidates who are both desirable to the employer, and likely to accept an offer. Recruiters must propose salary ranges that will attract good candidates, but without spending beyond necessary to secure acceptance from a quality candidate (Adler, 2020). Many subjects likely understood these goals based on past experience (100% had prior experience in this role). Because it is common to provide workers with some direction for a hired job, we stated them in the task instructions. The recruiter's incentives both inside our experiment and more broadly discourage offers that are too low (because candidates will reject low offers), and offers that are too high (because the higher amount will eat into company resources). Rather than mechanically rewarding lower or higher pay, the recruiter must balance the costs and benefits of yield with using more resources.

Recruiters were paid their posted hourly rate for the task. Recruiters face reputational incentives to perform well; happy customers could lead to repeat business, high ratings, or referrals. Because performance pay is common, we also included a bonus, which we described in simple, non-technical language outlining the goals above. Additional details

⁹Recruiters could choose a maximum between \$20,000 to \$200,000 in \$10,000 increments. One can think of this as asking, "Should the firm be willing to pay \$20,000, \$30,000 etc. up to \$200,000?" Recruiters did not appear to feel constrained by this range since the minimum valuation was \$60,000 and the maximum valuation was \$180,000.

appeared in a FAQ portion of the instructions. Those who read the full details saw that bonuses for each job candidate were paid only if the candidate is hired (i.e., passes employer screening and accepts the job offer). The amount of the bonus depends on the salary used to attract the candidate, and was slightly lower if a large salary was used to attract the candidate. The bonus includes a measure of candidate quality (assessed by the employer) to allow recruiters to pay more to attract a higher quality candidate without penalty. Because there were no actual candidates nor firms, the bonuses were based on simulated outcomes based on data from comparable settings.

3.3 Experimental Manipulation

Our field experiment contains experimental manipulation on both the employer side *and* the candidate side. Example applications which show some of these manipulations can be found in Appendix L Figures L.1-L.4.

Employer Side. We manipulated whether the employer’s job applications asked the candidates for their previous salary, or not. When it was present, the question asked for the applicant’s annual base salary at their current or most recent job.

Candidate Side. We randomized the candidate’s answers listed on the job form. Our candidate randomizations fall into two categories. The first is related to the candidate’s biographical details, and the second is related to the candidate’s salary disclosures.

Biographical Details. To randomize gender, we created candidate names using the top four male and the top four female names from American cohorts of 1991-1994 according to the Social Security Administration, making the candidates a few years out of college at the time of our experiment.¹⁰

Another important randomized biographical detail is about low- and high-wage current firms. To randomize the current firms of our candidates, we used a list of the 12 biggest employers of software engineers from Monster.com and Indeed ([Monster](#), [Rayome \(2017\)](#)). These lists include substantial variation in median salaries for software engineers who recently graduated. According to PayScale.com, the highest was Oracle (median salary of \$126K), and the lowest was General Dynamics (median salary of \$73K).

¹⁰The names were: Andrew, Christopher, Joshua, Tyler, Emily, Jessica, Samantha, and Sarah.

To present job applications to recruiters, we needed additional characteristics for candidates besides a first name and a former employer. Our goal with these other characteristics was to hold them roughly constant at values representative of the broader software market. Some details were held constant: all candidates held a bachelor’s degree in computer science, and none required a work visa. However, we permitted some additional random variation in other biographical details in order to create natural variation in candidates (Appendix F lists these details).

Using this procedure, we created 32 candidate biographies; a “biography” consists of the specification combination of name, gender, previous and current employers, job titles and descriptions, dates of employment, undergraduate education, and skills. Each recruiter was given a packet of eight candidates to evaluate with each packet containing four male and four female candidates, with randomly chosen former employers. Each packet was then assigned to a treatment and subtreatment condition, described below. Thus, these 32 candidates were evaluated under different experimental circumstances (one packet per recruiter). By asking recruiters to evaluate the same 32 candidates, our experiment permits estimation with “biography fixed effects”, that is we can hold candidate characteristics fixed and see how recruiters react to the same candidate in different experimental conditions.

Salary History Disclosures. Candidates’ salary history disclosures (or lack thereof) were also randomized. In packets where candidates were asked salary histories, those who disclosed answered the question on the form by entering a number on the line. In packets where candidates were *not* asked, candidates disclosed using an optional field for “Additional Skills and Information. This section also included information about programming skills the applicant had; For disclosure without a prompt, the candidate would add a sentence such as “Current Base Salary: \$X yearly.” We randomized whether the statement appeared before or after the additional skills, as well as the language of this statement. Real job candidates vary widely in whether they choose to disclose their salary history with our own survey finding that 52% of job seekers only disclose if asked, while 28% always disclose and the remaining 19% never disclose (Cowgill et al., 2022). Examples of how prompted and unprompted disclosures looked on the applications can be seen in Appendix Figures L.1-L.4.

The amount of disclosed salary is also randomized among those who disclosed. These disclosure amounts were consistent with the candidate’s prior employer and gender, but also included some random variation (conditional on the biography). For each candidate’s current employer, we looked up the distribution of salaries for software engineers at the candidate’s location and job level using Payscale.com.¹¹ Candidates were assigned a salary

¹¹Data from websites like Payscale.com and Glassdoor.com are self-reported by workers who visit these

that was either relatively high (near the 75th percentile) or relatively low (near the 25th percentile) within their current firm’s salary distribution. This creates a wage residual for each candidate that recruiters could not anticipate from the biographical details listed on the application.

Figure 1: Distribution of Disclosed Salary Amounts



Notes: This figure shows the distribution of male and female salaries (among the job applications containing a disclosure). The mean disclosed salary amongst those who disclosed was \$89,565 for female candidates and \$104,638.

To understand how disclosure affects wage gaps, we built a gender wage gap into the salaries disclosed by our candidates. Among all salaries reported across recruiters in our experiment, the average female disclosed salary was 86% of the average male disclosed salary — however the actual wage gaps recruiters observe within their eight applications varied.¹²

websites, which could mean these data are inaccurate. However, Glassdoor has periodically compared its data to that provided by the Census, and they’ve found that the distribution of base salaries reported are very similar (Glassdoor, 2019). We used data from Payscale.com because it offers the ability to see the distribution of salaries by company, job roles, city and level of experience (a level of granularity that is not publicly available from sources like the CPS, BLS, the Census or Glassdoor.com).

¹²Amongst recruiters who saw any disclosed salaries, the gender gap they saw varied from females disclosing salaries that were on average 61% lower to 24% higher (at the 1st and 99th percentiles of the gender gaps seen by recruiters). To create the gap, female candidates who work at the same current firm/location were given a salary that was 85% compared to the male salary at the same firm/location. Although the wage gap conditional on working at the same firm/location is 85%, the actual average gap in wages seen over all salaries disclosed was 86% because we did not show every single firm/location equally across the full set of recruiters. We analyzed data from the 2015 American Community Survey (ACS) on individuals in computing jobs and found women earned around 85% compared to men in this data source. Publicly available salary data about specific firms – including the sources we used above (Glassdoor and PayScale) and all others we consulted –

Figure 1 shows the resulting distribution of disclosed salaries amongst candidates who disclosed for both male and female candidates. Despite the overall average gender wage gap built among our candidates, most variation in the disclosed wages can be explained by previous employers. Appendix F contains additional details about how we assigned salaries to candidates.

3.4 Subtreatments

We randomized salary prompts and disclosures within packets received by recruiters in three main ways. First, we randomized whether there was a prompt for salary history, holding the other employer features constant within each pack of eight applications. Recruiters were hired to screen applications from one firm, and thus we kept the application materials consistent within the recruiter’s packet. All eight applications asked for salary history, or all eight did not.

Second, we randomized candidates’ disclosure choices. Each packet contained either zero, four or eight candidates disclosing. When four candidates disclosed, we randomized which four, and sent a separate packet to another recruiter flipping the candidates’ disclosures. This grouping allowed us to measure how candidates’ disclosure choices affected each other through spillovers.

Finally, we randomized candidates’ disclosure amounts. As previously mentioned, candidates’ disclosure amounts were randomized (for those who disclosed). We grouped these disclosure amounts so that in some cases, more disclosing candidates were on the high end of their previous employer. This allowed us to measure and control for potential spillover in the monetary amount disclosed (i.e., how one candidate’s high disclosure affects other candidates’ outcomes). For the full details of the randomization described above, see Appendix D.

Balance. Appendix F.1 and D.1 show our candidate and employer/prompting manipulations were uncorrelated with the characteristics of the assigned recruiter or with each other (by construction).

do not contain gender-specific wage values.

4 Empirical Specifications

Our conceptual framework shows how the effects of disclosure on hiring outcomes are ambiguous, both in direction and magnitude. To measure the directions and sizes of effects in our empirical setting, we mainly use the two specifications below. In these regressions, the outcome variables y_{ij} are the assessments given to candidate i from recruiter j (such as the callback choice, salary offer, WTP and other assessments).

4.1 Specification 1: Effects of Disclosing by Amount

$$y_{ij} = \beta_1 \text{Disclosed}_{ij} + \beta_2 \text{Disclosed}_{ij} \times \text{AmountDisclosedDemeaned}_{ij} + v_i + \gamma_j + \beta_3 [\text{SpilloverControls}_{ij}] + \beta_4 [\text{RecruiterControls}_j] + \epsilon_j \quad (4)$$

where i indexes candidates and j indexes recruiters. *RecruiterControls_j* includes the gender, race, experience level, and hourly rate of the assigned recruiter (all balanced by design, section D.1). v_i signifies candidate biography fixed effects (see section 3.3). We include sub-treatment fixed effects γ_j which control for whether a salary history prompt was included on all or none of the 8 applications seen by a recruiter, the number of salaries disclosed, and the exact mix of salaries disclosed in the packet. Spillover controls are detailed in Appendix G.¹³

The variable *SalaryDisclosed_{ij}* takes the value one if a candidate disclosed their salary, and zero otherwise. For those who disclose an amount we subtract off the mean of disclosed amounts within gender, \$89,565 for female candidates and \$104,638 for men to create the variable *AmountDisclosedDemeaned_{ij}*. We set *AmountDisclosedDemeaned* = 0 for candidates that did not disclose (their overall impacts are captured by “*SalaryDisclosed* = 0”). Demeaning allows the coefficient on *SalaryDisclosed_{ij}* to have an easier interpretation as detailed below.

Interpretation. The slope coefficient on *SalaryDisclosed_{ij}* \times *AmountDisclosedDemeaned_{ij}* (β_2) tells us how much outcome y_{ij} changes with each additional dollar (above the mean) of salary a candidate discloses. Larger coefficients on *SalaryDisclosed_{ij}* \times *AmountDisclosedDemeaned_{ij}* (steeper slopes) indicate greater impacts on a recruiter’s evaluation. Greater magnitudes indicate recruiters updating beliefs about the candidate more based on the disclosure amount.

¹³Our main interest in this paper is the average direct effects of one’s disclosures on one’s own outcomes, but as described in Section 3.4, our experiment was also designed to study potential spillovers between candidates’ disclosures. Our regressions thus include a set of *SpilloverControls_{ij}*. We treat these as control variables and do not report spillover coefficients in this paper.

The opposite extreme (a low, flat $SalaryDisclosed_{ij} \times AmountDisclosedDemeaned_{ij}$ coefficient) indicates recruiters do not change their assessments or incorporate new information from the disclosure amount much. We will use this specification to measure the monotonicity assumption in our theoretical framework (Assumption 2). This requires that salary history has a greater slope with valuation than outside options (or vice versa). A positive coefficient on $SalaryDisclosed_{ij}$ (β_1) means that outcome y_{ij} is larger when disclosing a salary at the mean of disclosed amount within gender as compared to not disclosing.

We extend this specification in three main ways: First, we introduce interactions with whether the candidate was asked for their previous salary by a prompt on the application form, $Prompt_j$. Interactions with $Prompt_j$ test for differences in the association between salary and outcomes by whether salary history was prompted on the application form. Second, we instead introduce interactions with whether the job candidate being evaluated is female, $Female_{ij}$. The coefficient on the interaction of $Female_{ij} \times Disclosed \times AmountDisclosedDemeaned_{ij}$ indicates whether an extra dollar disclosed is treated differently by whether that disclosure came from a male or female job candidate. While the interaction of just $Female_{ij} \times Disclosed$ indicates if men and women disclosing at the mean of disclosed amount within gender are treated differently. Third, we have a fully saturated model with interactions with both $Female_{ij}$ and $Prompt_j$ to test for variation in the way male and female candidates are evaluated with and without a salary history prompt.

Disclosure Indifference Salaries. We can also use the coefficients estimated from Equation 4 to impute salaries where (on average) candidates would be indifferent between disclosing and not (based on the linear structure of Equation 4). We describe this procedure in more detail in Appendix H, and summarize the basic intuition here. The salary history level h at which candidates would be (on average) indifferent about disclosing (or revealing h) is the one where the intercept effect of disclosing (β_1 in Equation 4) is fully counteracted by the slope effects coming from $\beta_2(h - \mu)$, where μ is the average disclosed wage used to de-mean the disclosed salaries.

We calculate this value separately for men and women using the procedure in Appendix H and the augmented version of Equation 4 including gender interactions. These values tell us the salaries below which the average candidate in either gender would be better off being silent, and above which the average candidate would be better off disclosing. We then compare these values to the distribution of female and male salaries in our experiment to understand the costs and benefits of disclosing at all versus silence.

4.2 Specification 2: Effects of Disclosing a High versus Low Salary

The coefficient on whether the candidate disclosed a salary in Equation 4 tells us something about silence, is it better to disclose a salary at the mean or not at all? To get a broader idea of how recruiters interpret silence, we also compare not disclosing to disclosing a relatively high or relatively low salary conditional on the firm that the individual works for. When we constructed candidate salaries (3.3), three factors went into candidate salaries: 1) Some work at higher or lower -wage firms, 2) some are male or female, and 3) some are relatively well- or poorly- paid within their firm's distribution. Each candidate presents a single number (the sum of these factors), not individual components. Nothing in our materials provides *any* explanation for how candidate's prior wages were set (much less the three reasons above).

Nonetheless, recruiters could anticipate. Some employers are known for paying well. Recruiters who see a candidate coming from Apple may adjust expectations. Similarly, recruiters who know the gender wage gap may adjust expectations of salary history downward for women. The only source of variation that *cannot* be anticipated by recruiters is the within-firm variation, which we designed to be uncorrelated with any observable feature. Our job candidates either disclose at the 25th or 75th percentile of the wage distribution within firm and gender. To assess the effect of disclosing a relatively higher or lower salary that cannot be attributed to the different firms or candidate gender we present a new model.

We now use the 25th and 75th percentile terms in Equation 4 instead of the continuous variable *AmountDisclosedDemeaned_{ij}* and include interactions with whether the job candidate being evaluated is female.

$$\begin{aligned} y_{ij} = & \beta_1 \text{Disclosed25th\%ileSalary}_{ij} + \beta_2 \text{Female}_{ij} * \text{Disclosed25th\%ileSalary}_{ij} + \\ & \beta_3 \text{Disclosed75th\%ileSalary}_{ij} + \beta_4 \text{Female}_{ij} * \text{Disclosed75th\%ileSalary}_{ij} + \quad (5) \\ & v_i + \gamma_j + \beta_5 [\text{RecruiterControls}_j] + \beta_6 [\text{SpilloverControls}_{ij}] + \epsilon_j \end{aligned}$$

Disclosed25th\%ileSalary_{ij} takes the value one if a person disclosed a salary at the 25th percentile within their firm within their gender, while *Disclosed75th\%ileSalary_{ij}* is analogous for those disclosing at the 75th percentile. Our candidates only disclose at the 25th or 75th percentile, so the omitted group is those who did not disclose a salary. *Female_{ij}* takes the value one if the candidate is female.

Interpretation. The *Disclosed25th\%ileSalary_{ij}* coefficient (β_1) measures whether those disclosing at the 25th percentile are (on average) rewarded more or less than silent male candidates. Likewise *Disclosed75th\%ileSalary_{ij}* (β_3) does the same for those men disclosing at

the 75th percentile. A positive coefficient on $Disclosed75th\%ileSalary_{ij}$ (β_3) indicates that recruiters believe positive self-selection into disclosing (i.e., above-average types choose to disclose). Interactions with $Female_{ij}$ explores whether there are differential effects for disclosures by male versus female candidates.

Table 1: Summary Statistics of Recruiter’s Candidate Assessments

	All			Male Candidate			Female Candidate		
	(1) All	(2) Salary Disclosed	(3) Not Disclosed	(4) All	(5) Salary Disclosed	(6) Not Disclosed	(7) All	(8) Disclosed	(9) Not Disclosed
WTP	107,101	109,922	102,978	110,814	114,873	104,882	103,388	104,970	101,075
Outside Option 5th %tile	88,836	93,772	81,622	92,664	99,727	82,342	85,007	87,816	80,902
Outside Option 50th %tile	100,044	103,635	94,796	104,235	109,783	96,127	95,852	97,486	93,465
Outside Option 95th %tile	112,161	113,546	110,136	117,517	119,697	114,332	106,804	107,395	105,940
Offer	100,957	103,993	96,521	104,588	109,107	97,983	97,327	98,879	95,058
Callback	0.633	0.628	0.641	0.632	0.613	0.659	0.635	0.643	0.623
Offer CB	94,789	99,809	87,200	99,027	104,658	89,707	90,517	94,558	84,932
Observations	2048	1216	832	1024	608	416	1024	608	416

Notes: Each of our 256 recruiters evaluated eight candidates for a total of 2048 candidate level observations. Outcome variables are defined in Section 3.2.

5 Results

Table 1 shows descriptive statistics for the recruiter responses from our experiment. We can see that WTP, recruiter perceptions of outside options, and salary offers are higher for those who disclose versus those who do not; and for male candidates versus female candidates. However, these responses are shaped by the amount disclosed in addition to the fact of disclosure, and cannot tell us about some of our motivating questions (e.g., whether the candidate would have been better off disclosing or not). Using the results of our field experiment and the estimation strategies described in Section 4, we now explore how recruiters interpret the existence of salary disclosures and the amount of salary disclosed by our candidates. We estimate whether these responses were different if the salary was disclosed (or withheld) in a response to a prompt or offered voluntarily. We also estimate how the recruiter’s interpretation of candidates’ disclosures differed by gender.

We first measure how recruiters react to higher or lower disclosed salaries in the odd columns of Table 2 (Equation 4). On average, higher salaries increase recruiters’ WTP, beliefs about outside options, and TIOLI salary offers. This is consistent with employers believing that prior salary carries information about worker quality. On average, every additional \$1 of current salary disclosed increases WTP by \$0.64. Column 3 shows that recruiters’ expectations about outside options similarly increase with higher amounts, for every \$1 of current

salary disclosed, this increases by \$0.73. And with these updates, they increase the amount of salary they offer by \$0.66 for every \$1 of salary disclosed.

A key question is whether recruiters interpret salary disclosures differently depending on whether job candidates *volunteered* unprompted, rather than provided disclosure in response to a question. Salary disclosures are interesting partly because bans have stopped employers from asking or prompting them; these bans, however, do not prevent voluntary disclosure. The even columns of Table 2 add in an interaction with whether there was a prompt on the application and show that whether this disclosure was in response to a prompt or an unprompted disclosure on the part of the candidate does not seem to matter for this slope, coefficients on the interaction between disclosure amount and prompts are small and statistically insignificant.

Table 2: Level of Salary Disclosed (Prompted and Unprompted)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer	Call-back	Call-back
Disclosed	0.52*** (0.13)	0.74*** (0.19)	0.61*** (0.13)	0.72*** (0.18)	0.58*** (0.12)	0.72*** (0.18)	-0.00 (0.05)	0.03 (0.06)
Disc x Amount (Demeaned)	0.64*** (0.05)	0.63*** (0.07)	0.73*** (0.05)	0.69*** (0.07)	0.66*** (0.05)	0.66*** (0.07)	-0.02 (0.02)	-0.04+ (0.02)
Disc x Prompt		-0.43 (0.26)		-0.22 (0.25)		-0.28 (0.23)		-0.06 (0.09)
Disc x Amount (DM) x Prompt		0.02 (0.09)		0.06 (0.08)		-0.01 (0.08)		0.04 (0.03)
Mean Non-Disclosers (Prompt): <i>All</i>	10.30		9.48		9.65		0.64	
Mean Non-Disclosers (No Prompt): <i>All</i>		10.27		9.35		9.61		0.64
R ²	0.38	0.38	0.52	0.52	0.43	0.43	0.04	0.04
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header. WTP is the maximum salary offer from the candidate the firm should accept. Outside Option is the amount the candidate would be likely to accept with a 50% probability (a.k.a. the 50th percentile of the distribution of the candidate's outside options; the 5th and 95th percentile are available in the Appendix). Offer is the take-it-or-leave-it (TIOLI) offer for the candidate. Callback is whether the candidate was recommended for an interview or not. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Amount Demeaned is the disclosed amount demeaned by subtracting the mean within gender of all the disclosed salaries (8.95 for female candidates and 10.46 for male in \$10K increments) from the amount disclosed by a specific candidate on their application form. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

The coefficient on whether a candidate discloses at all in these regressions implies that candidates who disclose at the mean of disclosed salaries are better off than candidates who do not disclose (e.g. their salary offers are \$5800 (6%) higher than non-disclosing candidates in column 5). Disclosing in response to a prompt seems to lower the benefits of disclosing at

the mean, though these coefficients are again not statistically significant.

The lack of prompt effects in our results do not mean that prompts do not matter. Instead, it suggests that the prompts matter via their impact on a candidate's choice to disclose (or not), rather than through the employer's interpretation of disclosures when they happen. Survey evidence suggests that prompting does have a large effect on candidates' choice to disclose, particularly for women (Cowgill et al., 2022).

5.1 By Gender

Do the recruiters' interpretation of salary disclosures differ by the gender of the candidate? Table 3 adds in interactions with gender to Equation 4.

For our main outcomes, disclosure amount slopes are similar by gender. Recruiters increase their willingness to pay and salary suggestions by similar amounts for each additional \$1 disclosed by a candidate regardless of their gender. This implies, conditional on disclosure, if pre-existing gender disparities exist they will continue to be perpetuated by disclosure. Similar to our overall effects, across most outcomes we find no statistically significant evidence that prompted (versus unprompted) disclosures change recruiter evaluations for men or women, though there is a marginally significant effect on outside options for women who disclose higher amounts when prompted.

The impact of disclosing a salary at the mean versus not disclosing at all is different by gender: Men who disclose a salary at the mean get salary offers \$9,400 (10%) higher than men who do not disclose, but for women it is only \$2,300 (2.5%, $p=0.08$) higher. As we discuss next, this suggests that the effects of being silent or not will be different for men and women. Whether a candidate is better off disclosing (or silent) will depend on the level of their salary and their outcome from disclosing (which we explore next). Table 3 includes the outcomes for all job candidates evaluated. For discussion of just those candidates who were chosen for callback or the top candidate see Appendix Tables K4 and K5.¹⁴

¹⁴One may also wonder about how results vary by whether our recruiters felt our candidates were reporting truthfully as described in Appendix E. A little over a third of our recruiters stated that they felt job candidates might still mis-report their salary even if it could be verified later, and they stated that on average they expected job candidates to state a salary which was 9% higher than their true salary. Our results are quite similar if we subset only on the two-thirds of recruiters who felt job candidates would not mis-report (Table K6).

Table 3: Gender, Disclosure and Amount

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer	Call- back	Call- back
Disclosed	0.83*** (0.16)	1.11*** (0.23)	1.08*** (0.15)	1.18*** (0.21)	0.94*** (0.14)	1.16*** (0.21)	-0.00 (0.05)	0.03 (0.06)
Female x Disclosed	-0.60*** (0.15)	-0.74*** (0.20)	-0.92*** (0.12)	-0.92*** (0.17)	-0.71*** (0.14)	-0.87*** (0.18)	0.07+ (0.04)	0.05 (0.06)
Disc x Amount (Demeaned)	0.67*** (0.06)	0.67*** (0.09)	0.78*** (0.05)	0.77*** (0.09)	0.69*** (0.06)	0.70*** (0.08)	-0.02 (0.02)	-0.04+ (0.02)
F x Disc x Amount (DM)	0.01 (0.05)	-0.02 (0.07)	-0.00 (0.04)	-0.08 (0.06)	0.02 (0.04)	0.01 (0.06)	0.02 (0.02)	0.04 (0.02)
Disc x Prompt		-0.57+ (0.31)		-0.23 (0.29)		-0.43 (0.28)		-0.06 (0.09)
F x Disc x Prompt		0.29 (0.28)		0.03 (0.24)		0.30 (0.26)		0.05 (0.08)
F x Prompt		-0.27 (0.20)		-0.17 (0.17)		-0.16 (0.19)		-0.07 (0.06)
Disc x Amount (DM) x Prompt		0.00 (0.11)		0.02 (0.10)		-0.01 (0.10)		0.04 (0.03)
F x Disc x Amount (DM) x Prompt		0.04 (0.07)		0.12+ (0.06)		0.01 (0.07)		-0.02 (0.03)
Mean Non-Disclosers (Prompt):								
All	10.30		9.48		9.65		0.64	
Male	10.49		9.61		9.80		0.66	
Female	10.11		9.35		9.51		0.62	
Mean Non-Disclosers (No Prompt):								
All		10.27		9.35		9.61		0.64
Male		10.40		9.44		9.72		0.64
Female		10.14		9.26		9.50		0.63
R ²	0.38	0.38	0.54	0.54	0.44	0.44	0.04	0.04
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header. WTP is the maximum salary offer from the candidate the firm should accept. Outside Option is the amount the candidate would be likely to accept with a 50% probability (a.k.a. the 50th percentile of the distribution of the candidate's outside options; the 5th and 95th percentile are available in the Appendix). Offer is the take-it-or-leave-it (TIOLI) offer for the candidate. Callback is whether the candidate was recommended for an interview or not. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Amount Demeaned is the disclosed amount demeaned by subtracting the mean within gender of all the disclosed salaries (8.95 for female candidates and 10.46 for male in \$10K increments) from the amount disclosed by a specific candidate on their application form. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

5.2 The Effects of Silence

Our results above show that disclosing a higher salary amount improves recruiters' valuations for a candidate, their impressions about outside options, and the level of salary suggested. However, the choice to disclose the salary at all depends on the level of that salary and how recruiters interpret a candidate's silence. How do recruiters interpret silence? How does this differ by gender?

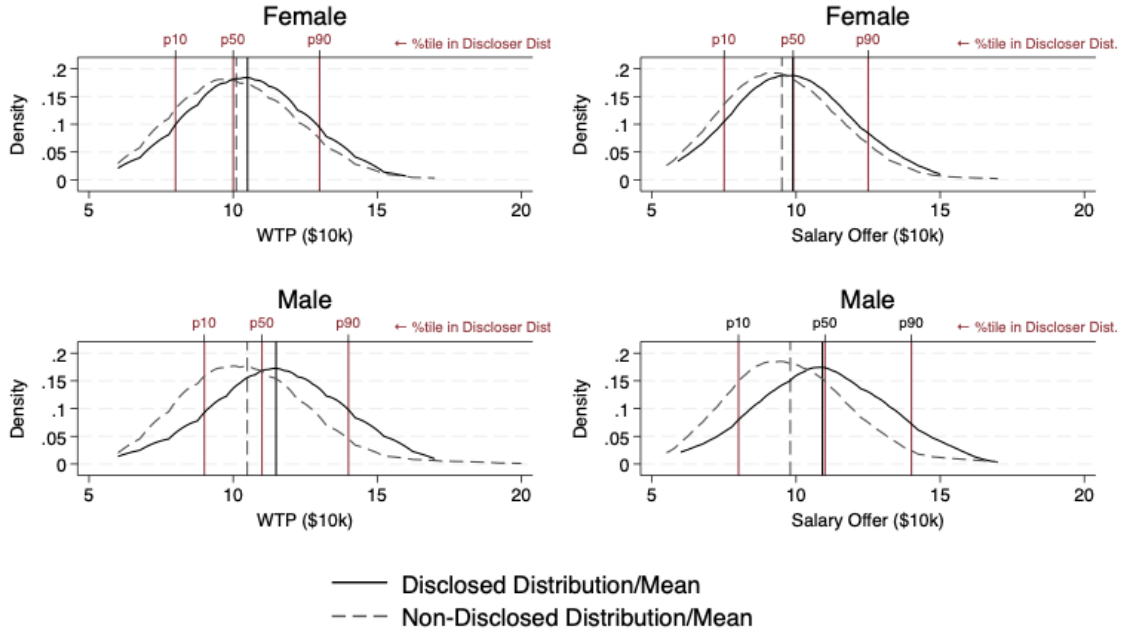
Our summary statistics in Table 1 imply that recruiters' WTP, beliefs about outside options, and TIOLI offers are lower for candidates who do not disclose (column 3) compared to those who do (column 2). In Figure 2, we illuminate the question of the cost (or value) of silence by gender. The figure shows the distribution of recruiter WTP and TIOLI salary offers for candidates who do not disclose versus those that do. We see that the WTP and Salary offers for those who do not disclose (dashed lines) are shifted to the left, and have lower means and quantiles. However, for female candidates the shift is relatively small: Female job candidates are punished *less* for non-disclosure than men.

On average, a silent woman is given a salary offer and WTP that is close to the 50th percentile of outcomes for all women who disclosed a salary. By disclosing her salary history, she moves the recruiter's offer and WTP up only slightly. This can be seen by how close the vertical dashed and black lines are in the Female panels of Figure 2. In contrast, the offers and WTP assessments for silent men are well below the 50th percentile of all male disclosers, while men who disclose their salary history are rewarded by recruiters with sizable jumps in offers and WTP. This can be seen in the large gaps between the vertical dashed and black lines for men in Figure 2. Women are punished less for staying silent than men. This pattern is qualitatively similar for recruiter's assessments of outside options (in Appendix Figure K1).

To delve into the effects of silence versus disclosing a relatively higher or lower salary by gender we turn to Table 4, where we examine the impact of disclosures by men and women at the 25th and 75th percentiles of their current employer's wage distribution as compared to those who do not disclose. When a man discloses at the 75th percentile, the recruiter's WTP, beliefs about outside option and TIOLI salary offer increase by about 12K-15K (12-16%) compared to not disclosing at all; that is, recruiters interpret silent men as significantly worse than men at the higher end of their within firm distribution. When a woman discloses at the 75th percentile, the increase is smaller in levels for outside option and TIOLI salary offer and represent 10-11% increases over non-disclosing women; recruiters interpret silent women as worse than 75th percentile women though not quite as much as they do for men.

In contrast, when a man discloses at the 25th percentile, we find no detectable benefit at

Figure 2: Distribution of Outcomes by Disclosure Status



Notes: This figure shows the distribution of recruiter’s choices. The solid and dashed density plot represents choices in response to a given salary disclosure for male and female candidates (among the job applications containing a disclosure). While the solid and dashed vertical lines represent mean outcomes for disclosers and non-disclosers respectively. The light gray vertical lines mark the percentile of the outcome in the disclosers’ distribution as listed.

all—statistically, they consider non-disclosing men to be similar to men at the 25th percentile of their within firm salary distribution. When a woman discloses at the 25th percentile, we find evidence of lower WTP, outside options and TIOLI salary (representing 2-3% declines in these outcomes) than if she had not disclosed at all (that is, non-disclosing women actually get slightly *better* outcomes than a woman disclosing at the 25th percentile).

Why would women be punished less for not disclosing? Our theory framework suggests that recruiters will punish silence less for candidates with higher disclosure costs. If a candidate faces higher disclosure costs, they may be silent even if their disclosure would be good (i.e., without the intention of hiding). In prior work, multiple studies suggest that women in particular do not like disclosing their salary (Goldfarb and Tucker, 2012; Sinha, 2020; Cowgill et al., 2022; Kuhn et al., 2022; Wang, 2023). These differences are robust to controlling for workers’s salaries, education and other characteristics; even well-paid women are reluctant to disclose. An alternative explanation for our results about silence is that recruiters underestimate the average market wages for this job. We explore this in Appendix

I, and infer that lack of knowledge does not seem to be driving our results.

Although recruiters penalize silent women less, their behavior has a flip side: The benefit of disclosing is also smaller for women (even among those who disclose relatively high salaries within their current firm).

Table 4: Average Effect of Disclosing a High versus Low Salary

	(1)	(2)	(3)	(4)
	WTP	Outside Option	Offer	Callback
Disclosed 25th %ile Salary	0.04 (0.15)	0.18 (0.14)	0.13 (0.13)	0.01 (0.05)
Female × Disclosed 25%tile Salary	-0.26+ (0.14)	-0.46*** (0.12)	-0.36** (0.13)	0.05 (0.05)
Disclosed 75th %ile Salary	1.23*** (0.18)	1.52*** (0.16)	1.35*** (0.16)	-0.01 (0.05)
Female × Disclosed 75%tile Salary	-0.22 (0.15)	-0.51*** (0.12)	-0.30* (0.14)	0.08+ (0.05)
Female 25th Disclosure Effect:				
<i>Total</i>	-0.22	-0.29	-0.23	0.06
<i>p-value</i>	0.16	0.06	0.11	0.22
Female 75th Disclosure Effect:				
<i>Total</i>	1.01	1.01	1.05	0.08
<i>p-value</i>	0.00	0.00	0.00	0.12
Mean Male Non-Disclosers	10.49	9.61	9.80	0.66
Mean Female Non-Disclosers	10.11	9.35	9.51	0.62
R ²	0.34	0.46	0.38	0.04
Observations	2048	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from Equation 5. Dependent variables are listed in the column header. WTP is the maximum salary offer from the candidate the firm should accept. Outside Option is the amount the candidate would be likely to accept with a 50% probability (a.k.a. the 50th percentile of the distribution of the candidate’s outside options; the 5th and 95th percentile are available in the Appendix). Offer is the take-it-or-leave-it (TIOLI) offer for the candidate. Callback is whether the candidate was recommended for an interview or not. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Disclosed Xth %tile Salary means a candidate disclosed a salary at the Xth percentile within their specific firm within their gender. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

5.3 Indifference about Silence

We can use the coefficients in Table 3 to estimate the thresholds for salary amounts above which men and women would be better off disclosing their salary and below which they would be better off withholding their current salary. We present these results in Table 5. Across several measures, men earning below $\approx \$91K$ are better off staying silent. This is the 28th percentile of the distribution of disclosed amounts in our experiment for men. Women who make below $\approx \$86K$ are better off staying silent. This is the 38th percentile of female

salary distributions. The lower percentile for men is consistent with our prior result that women are penalized less for staying silent.

Table 5: Disclosure Indifference Salaries (Implied by Table 3 Results)

	WTP	Outside Option	Offer
Male: Disclosure Indifference Salary	\$92,269 [2,291]	\$90,887 [1,833]	\$91,145 [1,940]
Female: Disclosure Indifference Salary	\$86,268 [2,204]	\$87,597 [1,787]	\$86,399 [1,913]
Percentile of Male Indifference	27.80	26.64	27.80
Percentile of Female Indifference	40.46	40.95	40.46

Notes: WTP is the maximum salary offer from the candidate the firm should accept. Outside Option is the amount the candidate would be likely to accept with a 50% probability (a.k.a. the 50th percentile of the distribution of the candidate’s outside options; the 5th and 95th percentile are available in the Appendix). Offer is the take-it-or-leave-it (TIOLI) offer for the candidate.

Recruiters act as if silent workers have a salary below the 28th (men) or 38th (women) percentiles. This is a significant discount. However, the theoretical literature on disclosure rationalizes even worse. Why aren’t recruiters more punitive? One possible reason is that salary histories can be “too high” as well as “too low.” As discussed previously, this is a key differentiator between salary history and other disclosure games. Our companion survey finds supporting evidence for this: workers are less willing to disclose extremely high salaries (?). Given this, silent candidates might be interpreted differently. Silent workers may contain a mixture of candidates whose salaries are too low, as well as some whose salaries are too high. This may partly explain why our recruiters do not assume that silent candidates are 0th percentile workers; in principle, some of the silent-types could be higher percentile workers hoping to avoid appearing overpriced.

5.4 Callback Decisions: WTP and Outside Options

Thus far we have discussed how salary disclosures impact recruiters’ willingness to pay, perceptions of outside options, and salary suggestions for our candidates. But in order for a candidate to benefit from a higher salary, the candidate needs to advance to the interview (and beyond). The choice to call a candidate back for interviews is another outcome in our experiment.

Theoretically, results about callbacks could go either way. If salary histories are very informative about outside options (and not about quality), disclosures could lower callbacks (especially for high amounts). The idea that a buyer would lower the quantity demanded in

response to a higher price appears in many economic frameworks (e.g., downward-sloping demand curves). In our own setup (Equation 2), this happens because higher prices diminish employer surplus from hiring a worker. Alternatively, it is also possible that salary disclosures communicate signals of quality more than outside options. In this case, a disclosure could increase callbacks.

Because we collected data about both quality (WTP) and outside options, we can estimate how these effects differ. On this point, our results suggest that disclosing an extra dollar of salary has a larger impact on recruiters' beliefs about outside options than on WTP. We see this in the coefficients on disclosure and amount for the outcomes of WTP and Outside Option in Table 2 (0.64 vs. 0.73 $p = 0.02$), and Table 3 (0.67 vs. 0.78, $p = 0.01$ for men and 0.68 vs. 0.78 $p = 0.01$ for women). Our results are more nuanced for the effect of disclosing at the mean, where the effects on Outside Option are clearly larger than those on WTP for men (Table 3, $0.83 < 1.08$ with $p = 0.01$). While the results are noisier for women (0.23 vs. 0.16, $p = 0.43$). Similarly in Table 4, disclosing at the 75th percentile has a greater effect on Outside Options than WTP for men ($1.23 < 1.52$, $p = 0.01$) but not for women (1.01 vs. 1.01, $p = 0.94$).

Given that disclosing often impacts Outside Options more than WTP, our framework suggests we should see lower callbacks for higher disclosed amounts. On this point our results are relatively noisy. In Tables 2, 3, and 4 do not find consistent effects of disclosing on callbacks.

We find some evidence that disclosing higher amounts unprompted lowers callback chances (Table 2). In Table 3, we find that higher unprompted disclosures by men lower callback chances (Table 3, Column 8). A few of our results suggest higher callback benefits for women disclosing compared to men disclosing at the same point in their respective distributions. Women are called back at higher rates when disclosing the mean for females (Table 3) compared to men disclosing at the mean for men. Table 4 shows that women disclosing at the 75th percentile for women have higher callback chances compared to men disclosing at the 75th percentile for men.

Despite this, our results about callbacks on the whole are noisy. Our estimates are underpowered because of the binary nature of the outcome variable. While we have some theoretical preconditions consistent with our model, our results on the callback outcome itself are not definitive.

6 Discussion

6.1 Generalizability of Results

Here we discuss the generalizability of our results. Our goal was to measure what information was conveyed by disclosure of voluntary signals in hiring, particularly signals correlated with gender or other protected categories that are at the center of recent bans. To make this goal tractable, we focused on salary history disclosures and gender in our setting. Our conceptual framework shows how results could differ in other settings. Appendix J discusses our findings using the SANS conditions (selection attrition, naturalness, and scalability) suggested by List (2020). We highlight a few key aspects of generalizability and external validity below.

Attrition. 17% of subjects did not complete the task. We find no evidence of differential attrition by observable characteristics of the recruiter. In addition, we find no patterns of differential attrition by the major treatment statuses (e.g. whether there was a prompt for salary history or the zero salaries disclosed, see Tables J1 and J2). We also tested for differential attrition based on the characteristics of the candidates who were included in the packet of job applications. We similarly see no patterns of differential attrition based on candidate characteristics, or their interaction with the prompt or whether a salary was disclosed (Table J3).

Scope of Work. Nearly all firms who use outsourced recruitment ask their recruiters to screen applications, and over 95% of recruiters have been asked to provide input about salary (Korn Ferry, 2018; SHRM, 2017; Staffing Industry Analysts, 2017). Even before the COVID-19 pandemic, over 80% of outsourced recruiting was performed remotely. Recruiters in this industry are often told to avoid searching on the Internet for information on job candidates, as this can violate employment law. For example the Equal Employment Opportunity Commission (EEOC) tells firms to avoid online searching for candidates (Claypoole, 2011). Our study design closely mimics each of these attributes.

Interviews. According to our survey, the most common way to ask salary history is on the job application (in writing). However, salary questions could also arise interactively in an interview. In this context, candidates and employers could exchange additional information to clarify the interpretation of the salary history. In a setting like ours, a candidate whose previous salary is “too high” could clarify their expectations and potentially avoid rejection.

Our experiment does not capture these effects. Because our experiment does not feature these clarifying questions, we isolate the effect of the salary history information and separate it from the disclosure of additional information (such as expectations or other mitigating circumstances). Communication of salary expectations is a separate, rich topic ([Roussille, forthcoming](#)). Some of our results suggest reasons why such clarifications would be useful: Salary histories alone could be either “too high” or “too low.” These additional, clarifying disclosures would provide mitigating effects for any policy implications within interview settings.

Recruiter Incentives. Bonuses — including both discretionary and formulaic kinds — are widespread in recruiting. Recruiters are primarily encouraged to fill open positions. This requires the recruiter to locate candidates who are both acceptable to the employer, and who are available to accept the job offer at mutually agreeable terms. To avoid wasting employer time on candidates unlikely to match, employers reward recruiters for selectivity.¹⁵

Because filling positions involves wage-setting, recruiters are rewarded for proposing wages that are acceptable to both the employer and the candidate. A comprehensive study of pay-setting practices by [Adler \(2020\)](#) shows that recruiters are tasked with “calibrating offers to maximize recruitment, with as little excess pay as possible.” The larger public is sometimes unaware of these incentives. In January 2022, an Atlanta-based recruiter bragged online about offering a candidate \$45K/year below the maximum allowed (the candidate accepted). She was heavily criticized in the media for underpaying the candidate ([CNN, 2022](#)). However, former and current recruiters commented that typical recruiting incentives reward low wage offers, as long as candidates will still accept them. A journalist described the incident as “ripp[ing] the band-aid off the dark and secretive world of salary negotiations” ([DiTomasso, 2022](#)). On her recruiting podcast ([Ruettimann, 2022](#)), HR consultant Laurie Ruettimann observed:

[W]hen I was making my way through HR, if the top range was \$130,000 and we paid someone \$85,000, I might get the difference as a bonus. People would be high-fiving me. [...] That was a way that HR and recruiting demonstrated cost savings. This is not something new[.] We were actually bonused for paying people below the midpoint.

At first glance, these incentives would appear to reward lower salaries mechanically. However, recruiters are also incentivized to make reasonable offers so that candidates will accept.

¹⁵According to industry reports, 60% of performance pay is measured by the number of positions filled ([SHRM](#)). In some cases, firms impose an explicit cap on the number of candidates a recruiter can forward. In settings where an employer is seeking all qualified candidates, firms may use an explicit monetary penalty for forwarding candidates who are a bad fit ([HR Magazine](#)).

Supply curves are upward-sloping. Higher pay is justified to recruit higher-quality candidates: “[T]he goal is to convince the desired candidate to accept the offer,” (Adler, 2020) but not “pay them above what’s necessary to secure their acceptance.”¹⁶ Textbooks for HR professionals suggest that recruiters, “control labor costs so that the organization’s prices of products or services can remain competitive” (Gerhart and Milkovich, 2019). Coverdill and Finlay (2017) emphasize that recruiters help the candidate find an acceptable offer, “without asking for too much.”

Rather than mechanically lowering salaries, recruiters must balance the costs and benefits of using more resources and changing the yield from offers. This tradeoff appears in many models of labor supply to an individual firm, starting as early as Hicks (1932) and Robinson (1933) as more recently summarized in Manning (2003) and Card (2022).

Setting & Conceptual Framework. Our empirical results most resemble Case 3 of our conceptual framework. They suggest that high h workers are better quality, but also have higher outside options (which diminishes the returns to hiring them). This is plausible in our setting. In the last decade, software firms have converged onto a class of common technologies for delivering applications from the cloud to devices. Star workers mastering these skills could have many competing offers. Those with fewer of these skills face idiosyncratic demand (depending on particular companies’ needs). In a different setting, rivals’ production technology may differ. “Cultural fit” could vary widely between firms. In these settings, rivals may disagree about which workers are “stars.” This could lead to Case 1 empirical results (“bargains” among high h), rather than those we find.

Using the language of List (2020), we intend this paper to be a wave 1 study. We focus on establishing causality and illuminating mechanisms based on a theory. Although our evidence comes from a particular industry (software engineering), our conceptual framework and research design can be adapted to other settings (and to adjacent research questions). Some results may be different in other settings. For example, we chose to have our fictitious candidates work in mid-size companies and possibly move to a start-up, this might have made our recruiters feel they needed to offer particularly high offers more so than a move to another mid-size firm.

We do not aim to generalize our results to the entire economy. To the contrary, our con-

¹⁶In some parts of the recruiting industry offered a bonus as a percentage of the candidate’s starting salary (Coverdill and Finlay, 2017). This may appear to give the recruiter a mechanical incentive to offer *higher* pay. However, recruiters must also make the offer acceptable to the employer, who could reject an excessively high offer. As with the case in the main text, the incentive to find mutually-acceptable terms offers some discouragement of extreme pay. Rather than mechanically increasing pay, the recruiter must balance the costs and benefits of using more resources with yield.

ceptual framework provides guidance about how our results depend on specific empirical properties and mediating factors of a given labor market (summarized in Section 2.3).

6.2 Implications for Bans

Salary disclosures are interesting partly because bans have been enacted to prevent employers from asking for salary history. Although our experiments are thematically related to bans, we do not attempt to make broad generalizations about their impacts. Our firm-side randomization (removing the question) does not necessarily generate the same effects as a ban. For example, a ban may bring about other punishment or stigma (real or imagined) for recruiters who use salary histories (even if they were voluntarily disclosed).

More generally, bans would generate long-term adjustment and general equilibrium effects as workers and firms react to the ban over time. Because of the relatively short nature of our experiment, our data does not capture these potential long-run impacts. A possible way to study this within our experiment might be to compare outcomes for recruiters in states where bans have been adopted to those where they have not (to see if recruiters in ban states behave differently). We examine this in Appendix Table K7, and cannot detect statistically significant effects of salary history disclosures based on the recruiter living in a ban state. Nonetheless, recruiters in these states may not have fully adjusted as they would over a longer horizon. Or, they may not have interpreted the presence or absence of salary history prompts as pertaining to bans.

In addition, any inferences from our experiment about the ban would still be sensitive to shifts in the share of people disclosing salaries (and related expectations). One possible implication from our results is that interpretation of non-disclosure can vary substantially depending on who discloses, and policies that shift disclosure rates could change this interpretation. Insofar as voluntary disclosure increases after the ban (possibly because of the issues in the disclosure and unraveling literature), these expectations could change on both sides of the market. However, because our candidates were fictitious, we do not know how they would respond to a salary history prompt or its removal in real life. Finally, our experiment is run in a specific setting of recruiting for software engineers. Even in this setting, we examine only one job type, and there may be limited scope for unobserved quality in this setting. Outcomes in other settings could differ, both in magnitudes and directions. Our conceptual framework even shows how the effects of different properties of a labor market could generate different results.

Rather than using our experiment to address the aggregate effects of bans in the long-term, we view results as generating hypotheses about effects that may (or may not) generalize, and

to highlight a few conceptual distinctions that might influence results. For example: The effects of the ban may vary depending on how employers use salary history (e.g., as a measure of quality or as a measure of outside option). This difference could influence whether bans affect not only salary levels, but also who is hired at all. In addition, salary histories could become revealed despite bans, through voluntary disclosure, inferences about silence, and the resulting unraveling. Finally, some of the negative effects of salary disclosure could be mitigated by employer/recruiter sophistication about disclosure behavior. For example, employers could anticipate aspects of disclosure behavior and/or amounts (such as the idea that female silence is less revealing). Separately, the disclosure of salaries is an interesting topic relating to wage-setting even without bans. Many people live in areas without a ban (or even a proposed ban), and the role of disclosures could be an important aspect of work in these settings.

7 Conclusion

What does a salary mean? In several classic models, workers are primarily differentiated by their outside options ([Manning, 2003](#); [Card et al., 2018](#); [Caldwell and Danieli, Forthcoming](#); [Jäger et al., forthcoming](#); [Lazear et al., 2016](#); [Burdett, 1978](#)), Salaries could reveal who is in a stronger bargaining position through their outside options. Another influential set of papers features workers differentiated by “hidden ability” ([Gibbons and Murphy, 1992](#); [Oyer and Schaefer, 2011](#); [Waldman, 2013](#); [DeVaro and Kauhanen, 2016](#)). Salary could be used as a signal of talent. A separate literature on voluntary disclosure attaches significance to saying nothing at all.

The informational content of salaries could affect any worker’s job search, or any employer’s hiring decision. However, the meaning of salary information is particularly relevant for job-seekers who are female, minority, or otherwise victims of discrimination. For such candidates, salary disclosures could perpetuate past disparities. Our research has aimed to bring these questions together. We began by developing a model that integrates the concepts above – signaling of ability, outside options, and voluntary disclosure.

To understand how salary disclosures impact recruiter assessments empirically, we executed a novel, two-sided field experiment in a real world setting. Our field experimental methodology allowed recruiters to be paid their normal hourly wage for assessing candidates, and allowed us to recover a richer set of assessments than the binary callback response that most audit studies rely on.

Through this field experiment, we find that higher salaries increase recruiter’s beliefs

about quality and outside options, and lead to higher salary offers for workers. We find no statistically significant differences in recruiters' assessments in prompted (versus unprompted) salary history disclosures, but prompts could still have effects if they compel workers to disclose.

Conditional on disclosing, we find no gender differences in the effects of disclosing extra dollars of additional current salary: each additional dollar disclosed nets similar increases in WTP, salary offers, and perceptions of the candidate's outside option for men and women. This does imply that salary disclosures could perpetuate past salary disparities for women and minorities who may have lower salaries due to discrimination or other structural reasons.

However, we do find gender differences in the effect of choosing to disclose at all (rather than staying silent). Men who disclose an average salary receive salary offers \$9,400 higher than men who do not disclose. For women it is only \$2,300 higher. These differences raise questions about the meaning of silence. We find that employers make negative inferences about silent candidates. Candidates without a disclosure are assumed to have lower-than-average quality and lower outside options, and are given lower salary offers. Recruiters are somewhat more lenient towards silent women. We find some evidence consistent with the idea that recruiters anticipate female discomfort about disclosure, and thus silence is less of a signal.

In theory, salary histories can be so high that it makes a worker unattractive. We have several findings consistent with this hypothesis. In particular, our results suggest that recruiters believe high-disclosing candidates deliver lower margins: The amount necessary to outbid their outside option is closer to their total value. This could lead to them receiving fewer callbacks. However, our empirical evidence for whether higher salaries actually decrease callback probabilities for candidates is noisy (partially due to the binary nature of the callback variable).

In this paper we have focused on salary history disclosure. However, our framework provides intuition about how results could plausibly vary for other disclosures, signals, and settings. We view our study as complementary with other research about how employers learn from voluntarily disclosed (or withheld) signals. A key part of our story is the imperfectly competitive labor market, where employers can use disclosure behavior to statistically discriminate about a candidate's outside options. Salary history can mean many things in many different models, particularly for groups whose salaries were historically depressed. Our paper aims to take a small step towards better understanding this topic.

References

- Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh**, “Perceived returns to job search,” *Labour Economics*, 2023, 80, 102307.
- Adler, Laura**, “What’s a Job Candidate Worth? Explaining Variation in Pay-Setting Practices,” *Working Paper*, 2020.
- Agan, Amanda and Sonja Starr**, “Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment,” *The Quarterly Journal of Economics*, 2018, 133 (1), 191–235.
- , **Bo Cowgill, and Laura Gee**, “Do Workers Comply with Salary History Bans? Voluntary Disclosure, Adverse Selection, and Unraveling,” in “AEA Papers and Proceedings,” Vol. 108 2020, pp. 33–37.
- Ashenfelter, Orley C, Henry Farber, and Michael R Ransom**, “Labor market monopsony,” *Journal of Labor Economics*, 2010, 28 (2), 203–210.
- Avivi, Hadar, Patrick Kline, Evan Rose, and Christopher Walters**, “Adaptive Correspondence Experiments,” in “AEA Papers and Proceedings,” Vol. 111 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2021, pp. 43–48.
- Barach, Moshe A and John J Horton**, “How do employers use compensation history? Evidence from a field experiment,” *Journal of Labor Economics*, 2021, 39 (1), 193–218.
- Bartik, Alexander and Scott Nelson**, “Deleting a signal: Evidence from pre-employment credit checks,” 2016.
- Bartos, Vojtech, Michael Bauer, Julie Chytilova, and Filip Matejka**, “Attention Discrimination,” *American Economic Review*, 2016, 106 (6), 1437–1475.
- Becker, Gordon M, Morris H DeGroot, and Jacob Marschak**, “Measuring utility by a single-response sequential method,” *Behavioral science*, 1964, 9 (3), 226–232.
- Bessen, James E, Chen Meng, and Erich Denk**, “Perpetuating Inequality: What Salary History Bans Reveal About Wages,” *Working Paper*, 2020.
- Biasi, Barbara and Heather Sarsons**, “Flexible wages, bargaining, and the gender gap,” *The Quarterly Journal of Economics*, 2022, 137 (1), 215–266.
- Black, Ines, Sharique Hasan, and Rembrand Koning**, “Hunting for talent: Firm-driven labor market search in the United States,” *Working paper*, 2020.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Beliefs about gender,” *American Economic Review*, 2019, 109 (3), 739–73.
- Burdett, Kenneth**, “A theory of employee job search and quit rates,” *The American Economic Review*, 1978, 68 (1), 212–220.

- Caldwell, Sydnee and Oren Danieli**, “Outside Options in the Labor Market,” *The Review of Economic Studies*, Forthcoming.
- Cappelli, Peter**, “Your approach to hiring is all wrong,” *Harvard Business Review*, 2019, 97 (3), 48–58.
- Card, David**, “Who Set Your Wage?,” *American Economic Review*, 2022, 112 (4), 1075–90.
- , **Ana Rute Cardoso, Joerg Heining, and Patrick Kline**, “Firms and labor market inequality: Evidence and some theory,” *Journal of Labor Economics*, 2018, 36 (S1), S13–S70.
- , **Fabrizio Colella, and Rafael Lalive**, “Gender Preferences in Job Vacancies and Workplace Gender Diversity,” Technical Report, National Bureau of Economic Research 2021.
- Claypoole, Theodore**, “EEOC regulations spotlight social media by Womble Bond Dickinson,” “<https://www.lexology.com/library/detail.aspx?g=1147c039-ef9c-4f6a-9ebb-448de20b8123>” 2011.
- CNN**, “Diversifying: Negotiating Your Way to Success,” “<https://www.cnn.com/audio/podcasts/diversifying/episodes/d19faee4-0af7-4047-b914-ae7001185387>” 2022.
- Corbae, Dean and Andrew Glover**, “Employer credit checks: Poverty traps versus matching efficiency,” *Review of Economic Studies*, forthcoming.
- Coverdill, James E and William Finlay**, “High tech and high touch,” in “High Tech and High Touch,” Cornell University Press, 2017.
- Cowgill, Bo, Amanda Agan, and Laura Gee**, “The Gender Disclosure Gap: Salary History Bans Unravel When Men Volunteer their Income,” *Working Paper*, 2022.
- **and Patryk Perkowski**, “Delegation in Hiring: Evidence from a Two-Sided Audit,” *Journal of Political Economy: Microeconomics*, forthcoming.
- Croson, Rachel and Uri Gneezy**, “Gender differences in preferences,” *Journal of Economic literature*, 2009, 47 (2), 448–74.
- Davis, Jesse, Paige Ouimet, and Xinxin Wang**, “Hidden Performance: Salary History Bans and the Gender Wage Gap,” *Working Paper*, 2020.
- DeVaro, Jed and Antti Kauhanen**, “An “opposing responses” test of classic versus market-based promotion tournaments,” *Journal of Labor Economics*, 2016, 34 (3), 747–779.
- DiTomasso, Alyssa**, “Viral Recruiter opens up about her explosive post,” “<https://www.audacy.com/star94atlanta/latest/star-exclusive-viral-recruiter-tells-all>” 2022.
- Doleac, Jennifer L. and Benjamin Hansen**, “The Unintended Consequences of Ban the Box,” *Journal of Labor Economics*, 2020, 38 (2), 321–374.
- Dranove, David and Ginger Zhe Jin**, “Quality disclosure and certification: Theory and practice,” *Journal of Economic Literature*, 2010, 48 (4), 935–63.

- Exley, Christine L and Judd B Kessler**, "The gender gap in self-promotion," *The Quarterly Journal of Economics*, 2022, 137 (3), 1345–1381.
- , **Muriel Niederle, and Lise Vesterlund**, "Knowing when to ask: The cost of leaning in," *Journal of Political Economy*, 2020, 128 (3), 816–854.
- Flory, Jeffrey A, Andreas Leibbrandt, and John A List**, "Do competitive workplaces deter female workers?," *The Review of Economic Studies*, 2015, 82 (1), 122–155.
- Friedberg, Leora, Richard M Hynes, and Nathaniel Pattison**, "Who Benefits from Bans on Employers' Credit Checks?," *The Journal of Law and Economics*, 2021, 64 (4), 675–703.
- Gee, Laura K**, "The more you know: Information effects on job application rates in a large field experiment," *Management Science*, 2019, 65 (5), 2077–2094.
- Gerhart, Barry and George T Milkovich**, *Compensation*, McGraw-Hill/Irwin Nova Iorque, 2019.
- Gibbons, Robert and Kevin J Murphy**, "Optimal incentive contracts in the presence of career concerns: Theory and evidence," *Journal of political Economy*, 1992, 100 (3), 468–505.
- Glassdoor**, "Progress on the Gender Pay Gap: 2019," *White Paper*, 2019.
- Goldfarb, Avi and Catherine Tucker**, "Shifts in privacy concerns," *American Economic Review*, 2012, 102 (3), 349–53.
- Grossman, Sanford J**, "The informational role of warranties and private disclosure about product quality," *The Journal of Law and Economics*, 1981, 24 (3), 461–483.
- **and Oliver D Hart**, "Disclosure laws and takeover bids," *The Journal of Finance*, 1980, 35 (2), 323–334.
- Hall, Robert E and Alan B Krueger**, "Evidence on the incidence of wage posting, wage bargaining, and on-the-job search," *AEJ: Macroeconomics*, 2012, 4 (4), 56–67.
- Hansen, Benjamin and Drew McNichols**, "Information and the Persistence of the Gender Wage Gap," *National Bureau of Economic Research Working Paper*, 2020.
- Hicks, JR**, "Theory of wages," 1932.
- Horton, John J**, "Price floors and employer preferences: Evidence from a minimum wage experiment," *American Economic Review*, Conditionally Accepted.
- HR Magazine**, "<https://www.shrm.org/hr-today/news/hr-magazine/pages/1103hirschman.aspx>".
- HRDive.com**, "Salary history bans: A running list of states and localities that have outlawed pay history questions," "<https://www.hrdive.com/news/salary-history-ban-states-list/516662/>".

- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer**, “Worker beliefs about outside options,” *Quarterly Journal of Economics*, forthcoming.
- Jin, Ginger Zhe, Michael Luca, and Daniel Martin**, “Is no news (perceived as) bad news? An experimental investigation of information disclosure,” *American Economic Journal: Microeconomics*, 2021, 13 (2), 141–73.
- Jobvite**, “Recruiter Nation Survey 2015,” https://www.jobvite.com/wp-content/uploads/2015/09/jobvite_recruiter_nation_2015.pdf 2015.
- Kessler, Judd B, Corinne Low, and Colin D Sullivan**, “Incentivized resume rating: Eliciting employer preferences without deception,” *American Economic Review*, 2019, 109 (11), 3713–3744.
- Khanna, Shantanu**, “Salary History Bans and Wage Bargaining: Experimental Evidence,” *Labour Economics*, 2020, p. 101853.
- Korn Ferry**, “Measuring Up? A new research report about RPO metrics,” <https://www.kornferry.com/insights/this-week-in-leadership/measuring-up-a-new-research-report-about-rpo-metrics> 2018.
- Kuhn, Peter and Kailing Shen**, “What Happens When Employers Can No Longer Discriminate in Job Ads?,” *American Economic Review*, 2023, 113 (4), 1013–1048.
- , **Shuo Zhang, Taoxiong Liu, and Kebin Dai**, “Should I Show or Should I Hide: When Do Jobseekers Reveal Their Wages?,” *Working paper*, 2022.
- Landay, Karen and Sarah DeArmond**, “Recruitment Process Outsourcing and Recruiter and Hiring Firm Characteristics,” *Journal of Personnel Psychology*, 2018.
- Laschever, Sara and Linda Babcock**, *Women Don’t Ask: Negotiation and the Gender Divide*, Princeton University Press, 2003.
- Lazear, Edward P, Kathryn L Shaw, and Christopher Stanton**, “Making do with less: working harder during recessions,” *Journal of Labor Economics*, 2016, 34 (S1), S333–S360.
- List, John A**, “Non est disputandum de generalizability? A glimpse into the external validity trial,” Technical Report, NBER 2020.
- Manning, Alan**, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2003.
- Marianne, Bertrand**, “New perspectives on gender,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 1543–1590.
- Mask, Joshua**, “Salary history bans and healing scars from past recessions,” *Labour Economics*, 2023, 84, 102408.
- Mathios, Alan D**, “The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market,” *The Journal of Law and Economics*, 2000, 43 (2), 651–678.

- Milgrom, Paul R**, “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 1981, pp. 380–391.
- Monster**, “Monster 2018 State of Recruiting Survey,” 2018.
- Monster**, “Companies Hiring Tech Workers Now,” “<https://www.monster.com/career-advice/article/companies-hiring-tech-workers-now>”.
- Murciano-Goroff, Raviv**, “Missing women in tech: The labor market for highly skilled software engineers,” *Management Science*, 2021.
- Niederle, Muriel and Lise Vesterlund**, “Do women shy away from competition? Do men compete too much?,” *The quarterly journal of economics*, 2007, 122 (3), 1067–1101.
- NPR**, “Planet Money: LeBron James Is Still Underpaid,” “<https://www.npr.org/transcripts/628132840>” 2018.
- Oyer, Paul and Scott Schaefer**, “Personnel Economics: Hiring and Incentives,” *Handbook of Labor Economics*, 2011, 4, 1769–1823.
- Pallais, Amanda**, “Inefficient hiring in entry-level labor markets,” *American Economic Review*, 2014, 104 (11), 3565–99.
- Rayome, Alison DeNisco**, “The 10 companies hiring more software engineers than anyone else in Silicon Valley,” 2017.
- Robinson, Joan**, *The economics of imperfect competition*, Springer, 1933.
- Roussille, Nina**, “The central role of the ask gap in gender pay inequality,” *Quarterly Journal of Economics*, forthcoming.
- Ruettimann, Laurie**, “Recruiter Mercedes Johnson Tells Her Side of the Story,” “<https://laurieruettimann.com/recruiter-mercedes-johnson-story/>” 2022.
- Sherman, Eliot L, Raina Brands, and Gillian Ku**, “Dropping Anchor: A Field Experiment Assessing a Salary History Ban with Archival Replication,” *Management Science*, 2023, 69 (5), 2919–2932.
- SHRM**, “Society for Human Resources Management: Talent Acquisition Benchmarking Report,” “<https://www.shrm.org/ResourcesAndTools/business-solutions/Documents/Talent-Acquisition-Report-All-Industries-All-FTEs.pdf>” 2017.
- SHRM**, “Society for Human Resources Management: Rewarding Recruiters for Performance,” “<https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/rewarding-recruiters-for-performance.aspx>”.
- Sinha, Sourav**, “Salary History Ban: Gender Pay Gap and Spillover Effects,” *Working Paper*, 2020.
- Sran, Gurpal, Felix Vetter, and Matthew Walsh**, “Employer Responses to Pay History Inquiry Bans,” *Working Paper*, 2020.

Staffing Industry Analysts, “RPO Market Developments,” 2017.

Stern, Scott, “Do scientists pay to be scientists?,” *Management science*, 2004, 50 (6), 835–853.

Viscusi, W Kip, “A note on” lemons” markets with quality certification,” *The Bell Journal of Economics*, 1978, pp. 277–279.

Waldman, Michael, “Classic promotion tournaments versus market-based tournaments,” *International Journal of Industrial Organization*, 2013, 31 (3), 198–210.

Wang, Shiya, “On the Basis of Leverage: What the Salary History Ban Reveals About How Employers Set Initial Salary Offers,” *Working Paper*, 2023.

Appendix: For Online Publication Only

A Labor Market: Software Engineering

We tasked our recruiting workforce with screening applicants for a software engineer position. The software sector is an ideal labor market for studying the effects of salary disclosures and bans on asking for salary history. The market for software jobs features several particularly attractive features for this study.¹

First, technical roles exhibit persistent gender wage and employment differences that span multiple decades (Blau et al., 2013; Goldin et al., 2017). Only 19% of computer science degrees are held by women, and one-third of workers in the technology sector in Silicon Valley are female.² Given the high wage and employment growth in this sector, technology may be a growing source of income inequality overall (Krueger, 1993; Acemoglu and Autor, 2011). Second, the technology sector features well-documented labor shortages and high levels of competition between employers for qualified workers. Technology executives regularly lobby Congress for expansions to the H1-B visa program to address the undersupply of software developers. Firms in this sector are generally interested in hiring multiple qualified candidates whenever possible. As the H1-B lobbying shows, hiring is limited not by demand, but by the supply of qualified workers. As a result, we can measure how salary disclosures and bans of prompting disclosures affect the likelihood of a candidate being called back, in addition to how the composition of selected candidates and salaries changes.

Second, by choosing this industry, we bias our study toward finding smaller differences between experimental variations. Labor shortages should erode the effect of gender and past salary on the evaluation of our job candidates. With strong competition for qualified candidates, there is likely to be less taste-based discrimination (Becker, 1957). This might lead salary history bans to be less effective in this industry. Behavioral economics phenomena such as “framing” and “anchoring” are often used to motivate why salary disclosures can be harmful and why salary history bans might reduce wage gaps. Effects in other, less-competitive sectors may be stronger.

B Details of Recruiter Selection

The recruiters in our experiment appeared on LinkedIn and UpWork offering recruiting services (both freelance and full-time), and we directed them through UpWork for the experiment’s payroll needs. Upwork allowed workers to either be paid an hourly rate or to negotiate a pay-by-task contract. Each recruiter’s profile includes an hourly rate suggested by the recruiter. We offered to pay our subjects the hourly rate posted on their profile. We

¹We chose to examine the market for engineers with moderate experience so that our candidates had a previous wage history that could (or could not) be disclosed.

²See https://nces.ed.gov/programs/digest/d18/tables/dt18_325.35.asp and <https://www.bloomberg.com/news/articles/2019-02-13/silicon-valley-is-using-trade-secrets-to-hide-its-race-problem>

also offered a bonus contract designed to align their interests with the firm's as they made decisions as these are common in recruiting. All recruiters worked remotely and corresponded with us directly over the Internet. Each qualified recruiter was sent the materials containing a set of applications to review and an online form to enter their assessments of each candidate. Recruiters were also sent a description of the firm and the hiring needs for the opening.

Each recruiter was required to sign a nondisclosure agreement, a common practice in real-world recruitment outsourcing in order to protect firm and candidate confidentiality. All these materials are available in the Section L. We did not directly tell recruiters that they were part of a larger recruiting workforce containing peers, but our instructions did reference the firm's other HR staff. The NDA also helped to address the possibility that recruiters would discover each other through circumstance and discuss the assignment. All recruiters signed the NDA, although some felt it was unnecessary because it was covered by the platform's terms of service.

To be eligible for an invitation into our workforce, recruiters on the platform had to be listed as independent (rather than affiliated with an agency),³ based in the United States according to their profile⁴ and had to have worked previously in real-world recruiting roles for office jobs.

We searched on keywords such as: "recruiter," "sourcing," "talent acquisition," "staffing," and "human resources." We did this in two waves. Wave 1 took place in the summer of 2018, while wave 2 was executed in late 2019. Over both waves, a list of approximately 20,000 possible recruiters was identified on keywords, then we examined a random sample of approximately 5,000 possible recruiters, and research assistants marked about 1,750 recruiters as qualified, by checking the recruiter's profile for prior real-world experience in hiring or recruiting for non-manual work. We then invited each qualified recruiter charging less than or equal to \$100 per hour.⁵ Approximately 400 wrote back in response to our inquiry to accept the job offer within the timeframe of our experiment. Most of the remainder did not write back at all; or write back after the experiment was completed. Some of these 400 were included in another study, and as such, are not reported on in this paper. We report on 256 recruiters who were part of this study.

These job requirements are typical for recruiting. The BLS's occupational data suggest that human resource work is mid-skill, work requiring a bachelor's degree, but no related work experience or prior on-the-job training.⁶ According to the BLS, our requirement of prior experience for recruiters is actually more stringent than a typical requisition for a recruiter.

³We did not hire agencies in order to avoid the possibility of recruiters in different treatment arms having discussions among each other.

⁴We focused on U.S. based recruiters who would be familiar with the qualifications of U.S. based candidates.

⁵The recruiters all indicated an interest in HR or hiring through the keywords they put in their profile. We also asked each invited recruiter for a résumé or LinkedIn profile. Before officially having them start the project, we checked these résumés or profiles for hiring experience. If the experience wasn't clear, we offered them the chance to clarify by asking them to tell us about their hiring experience. If this answer implied that a firm would be interested in hiring this person for this role based on their response, then we proceeded. Approximately 40 individuals who responded to our initial inquiry were ultimately not sent experimental materials, mostly because they had insufficient experience with hiring/recruiting/screening.

⁶<https://www.bls.gov/ooh/business-and-financial/human-resources-specialists.htm#tab-1>

Over 70% of our subjects reported over three years’ prior experience, and 98% stated that they provided salary input in prior recruiting assignments. We did not require prior experience specifically in recruiting software engineers. However, prior experience in software-engineering recruiting is not necessary for a recruiting job at many tech companies, as hiring for high-skills jobs is quite similar across many sectors (Adler, 2020).

Recruiters’ hourly rates were paid shortly after we verified their input was complete. Bonuses were paid between 30 to 45 days later.⁷

C Our Recruiters and BLS Averages

According to the BLS, in 2018 Human Resource workers across all industries were 69.7% female, 10.5% black, while the median hourly wage was \$29.01 across all industries, and \$41.93 in the software industry.⁸ As compared with the BLS statistics about human resource workers in the U.S., the recruiters in our study were slightly more likely to be female (75%), twice as likely to be black (23%), and had a higher hourly wage of \$44 (Table D1).

The BLS does not report demographic characteristics of industry \times occupation cells. However, these figures can be calculated using the Five-Year (2012-2017) American Community Survey Public Use Microdata. There are approximately 115 human resource specialists in the software industry in this sample. They are approximately 80% female, and 75% white.

D Details of Randomization Procedure

Our randomization procedure was sequential, proceeded in batches, and was designed to address covariate balance through re-randomization. For recruiters who were invited, accepted, and met our pre-screening qualifications (signed a non-disclosure agreement and possessed relevant experience), the recruiters’ demographics were manually coded.⁹ We merged the coded demographics data with data about the recruiter’s prior work experiences and posted wage rate.

Before sending out the experimental materials for recruiters’ feedback, we performed a covariate balance check (described below). If our covariate balance test passed, we would send the experimental materials to the recruiters. If the balance checks failed, we re-randomized the current batch (previous batches had already been sent to recruiters, who had already begun work on them, so they could not be re-randomized).

Our balance test checked for equality of the average of the following covariates across

⁷Because there were no actual candidates nor firms, they were based on simulated outcomes based on data from comparable settings.

⁸See <https://www.bls.gov/ooh/business-and-financial/human-resources-specialists.htm> and https://www.bls.gov/oes/current/naics4_511200.htm.

⁹For our full sample of recruiters, the recruiter’s self-reported gender matched our manually coded gender 99% of the time. The recruiter’s self-report of identifying as black matched our manual coding of this variable 92% of the time, while a recruiter’s self-report of identifying as white matched our manual coding 87% of the time.

treatment arms. The covariates were: 1) race (dummy variables for white and black), 2) gender, 3) the recruiters' advertised hourly rate, and 4) a dummy variable for whether the recruiter had previously logged hours on the website we used to hire them.

We tested for equality of these means across all treatment groups (a single test per variable for equality across all treatment arms). In addition, we tested for pairwise equality across all treatment arms. For assignments where these tests' p -values were less than 0.2, we re-randomized. We also randomized if the pairwise comparison for any two sub treatments was less than 0.05.

The sequential balance checks were cumulative. The tests above included observations for all prior assignments including the current batch. However, the current batch was the only batch that could be potentially adjusted if re-randomization was necessary. Batches were processed approximately once per week, so that recruiters would not have to wait long after accepting our offer to begin work.

Subtreatments. As described in Section 3.4, we randomized in three main ways: prompt included on job applications, candidate's disclosure choices, candidate disclosure amounts. The randomization produced 22 sub treatments, where a sub treatment is a combination of $\{\text{asked, not}\} \times \{\text{all disclose, half disclose, other half disclose, none disclose}\} \times \{\text{all high amounts, half high + half low, other half high + half low, all low amounts}\}$. Our total number of treatments is less than $2 \times 4 \times 4 = 32$ because in cases where no candidates disclose, amounts are irrelevant.

D.1 Recruiter Characteristics Balance

Our study randomized the salary history prompt, proportion disclosing, and distribution of amounts disclosed at the recruiter level. Prior research suggests that hiring decisions differ according to managers' characteristics.¹⁰ As such, we implemented a randomization procedure to guarantee covariate balance on recruiter characteristics such as race and gender across recruiter-level variations. This effectively implemented stratified randomization, guaranteeing that (for example) male recruiters were not over-assigned to one particular experimental arm by accident.

Table D1 shows that our stratification procedure succeeded; the recruiter demographics are balanced across whether the recruiter was shown applications with a prompt or not and whether the recruiter was shown zero, four, or eight candidates who disclosed. Almost none of the mean differences between our main experimental variations approach traditional levels of statistical significance.¹¹

¹⁰For example, [Giuliano et al. \(2009\)](#) report that nonblack managers hire more white workers and fewer black workers. [Dee, 2005](#) find that educators evaluate students of the opposite gender more harshly.

¹¹Proportion of screeners who are black is 28% for those shown four disclosures while it is 20% for those shown zero disclosures, a comparison which has a one-sided t-test of $Pr(T > t) = 0.0956$. The proportion who had been asked for salary input before is 100% for those shown zero disclosures while it is 97% for those shown four disclosures, a comparison which has a one-sided t-test of $Pr(T > t) = 0.0919$, a difference that is statistically significant but likely not economically significant. We randomized three things at the recruiter level: 1) prompt, 2) proportion disclosed, and 3) distribution of amounts disclosed. The interaction of those

Table D1: Recruiter Balance

	Female Re- cruiter	White	Black	3+ Yrs Exp	Hourly Rate	Asked Salary Input	% of Sam- ple
All Recruiters	0.75	0.52	0.23	0.71	44.07	0.98	100.0
No Salary Prompt	0.76	0.56	0.22	0.67	43.65	0.97	43.8
Has Salary Prompt	0.74	0.49	0.24	0.74	44.40	0.99	56.3
0 Salaries Disclosed	0.77	0.55	0.20	0.68	43.07	1.00	21.9
4 Salaries Disclosed	0.75	0.51	0.28	0.71	44.09	0.97	37.5
8 Salaries Disclosed	0.74	0.52	0.20	0.72	44.59	0.98	40.6
NoPrmpt 0Disc	0.72	0.56	0.19	0.66	44.97	1.00	12.5
NoPrmpt 4Disc MoreHigh	0.81	0.38	0.31	0.69	48.84	0.94	6.3
NoPrmpt 4Disc MoreLow	0.63	0.50	0.25	0.63	41.23	0.94	6.3
NoPrmpt 4Disc Mixed	0.88	0.69	0.31	0.69	37.32	1.00	6.3
NoPrmpt 8Disc AllHigh	0.75	0.63	0.13	0.63	41.88	0.88	3.1
NoPrmpt 8Disc AllLow	0.63	0.75	0.00	0.38	42.31	1.00	3.1
NoPrmpt 8Disc Mixed	0.88	0.56	0.25	0.88	46.13	1.00	6.3
Prmpt 0Disc	0.83	0.54	0.21	0.71	40.54	1.00	9.4
Prmpt 4Disc MoreHigh	0.81	0.56	0.13	0.75	42.88	1.00	6.3
Prmpt 4Disc MoreLow	0.63	0.50	0.44	0.75	46.38	0.94	6.3
Prmpt 4Disc Mixed	0.75	0.44	0.25	0.75	47.92	1.00	6.3
Prmpt 8Disc AllHigh	0.69	0.31	0.44	0.69	46.75	1.00	6.3
Prmpt 8Disc AllLow	0.56	0.50	0.06	0.75	40.06	0.94	6.3
Prmpt 8Disc Mixed	0.80	0.53	0.20	0.75	45.92	1.00	15.6

Notes: This table shows a subset of the demographics of our recruiting workforce of 256 recruiters by whether they were shown applications with a salary history prompt or not, whether they saw 0, 4, or 8 candidates disclose a salary, and by combinations of prompt/no prompt, 0/4/8 salary disclosures, and distributions of amounts disclosed. Columns are attributes of our recruiting workforce. 3+ Yrs Exp is an indicator for whether the recruiter self-reported they had worked in this type of work for at least 3 years. Hourly rate is the rate the recruiter had set on the platform for their hourly wage. Asked Salary Input is an indicator for whether the recruiter self-reported they had given salary input as part of this type of work in the past.

E Additional Questions

In section 3.2 we describe the bulk of our recruiting task. Here we detail some additional questions we asked.

Single Offer. We asked recruiters to tell us which of the eight candidates they would like to make an offer to if they could only make one single offer.

Number and Sources of Competing Offers. Recruiters were also asked to estimate how many competing offers each candidate would receive during his or her search from other employers. To simplify this task, recruiters could choose either “zero or one,” or “two or more.”¹² Recruiters also state whether competing job offers would come from the candidate’s own search efforts, or from rival employers’ search efforts.

Mis-reporting. We asked recruiters to consider job candidates like the ones they are reviewing for us when answering if these candidates would mis-report their most recent salary (even if the salary could be verified later). We then additionally asked the recruiters what they thought the true salary would be for a candidate who mis-reports a salary of \$90,000.

Demographics and Experience. We also asked our recruiters how long they had been doing this type of work and how often they were asked for salary input. We asked them to indicate which of the following they used to make judgments on salary related questions: used previous experience, looked up salaries on websites like glassdoor.com, spoke with others, and/or “Other.” We also asked them to self-report their gender and ethnicity.

F Details of Creating Candidates

For first names, we used the top four male and female names given to Americans according to the Social Security Administration (making job candidates between 24-27 years old at the time we began our experiment).¹³ We blacked out the last name so that recruiters could not try to contact our candidates or look them up online (Acquisti and Fong, 2015); we

three variations results in 22 distinct recruiter-level sub-treatments. In Table D1 we show the mean of the recruiter characteristics across these sub-treatments. There are a total of 546 two-way comparisons, and of these, 16% are statistically significant at traditional levels. As such, we include controls for screen characteristics in our models.

¹²Prior research suggests that outside options increase the bargaining power of the candidate (Blackaby et al., 2005), and that employed workers rarely receive more than one job offer at a time when searching. Faberman et al. (2022) find that only 29.1% of employed workers who are looking for work receive at least one offer per month.

¹³Male names were Andrew, Tyler, Joshua, and Christopher. Female names were Jessica, Emily, Samantha, and Sarah. See <https://www.ssa.gov/oact/babynames/top5names.html>. We excluded the name “Ashley” as it could be interpreted as being either male or female.

also encouraged recruiters to make decisions based on the application materials rather than investigating them online.

Each candidate was assigned a bachelor’s degree in computer science from universities ranked third to ninth in the country in computer engineering by *U.S. News and World Report*.¹⁴ We excluded the top two universities (MIT and Berkeley) to avoid the possibility that the top institutions might have some special cache, since variation in school quality was not one of the primary variations of interest for the experiment.

Previous firms were chosen from the top firms that hire software engineers.¹⁵ To ascertain previous salaries, we matched these firms with salaries reported on Payscale.com.¹⁶ Payscale.com provides very granular data indexed by company, job roles, city, and level of experience. We obtained the 25th, 50th and 75th percentile of salaries for software engineers with one to three years of experience in each firm’s headquarter cities.¹⁷

Each candidate’s biography required a realistic salary that could be disclosed when assigned to disclosure treatments. To approximate realistic gender gaps in salaries, we analyzed data from the 2015 American Community Survey (ACS).¹⁸

Our goal is to adjust the firm-city specific salaries from Payscale.com to create plausible male and female salaries for all candidate biographies. We adjust the Payscale.com salaries for men at each firm by multiplying the appropriate salary by 1.05. Then, we multiply the result by 0.80 to get the estimated female salaries at the same firm, location, and job. We derived these estimates from our analysis of the ACS data.¹⁹

The salaries reported on our job applications use these numbers, with a few additional

¹⁴There are in fact nine schools ranked between 3-9 as a result of ties. They are: Carnegie Mellon, University of Illinois Urbana-Champaign (UIUC), Georgia Tech, University of Michigan, University of Texas at Austin, Cornell University, Cal Tech, the University of Washington, and Purdue University. We randomly selected from the three schools tied for ninth place so that our final applicants did not attend Purdue University. See: <https://www.usnews.com/best-graduate-schools/top-engineering-schools/computer-engineering-rankings>.

¹⁵See <https://www.techrepublic.com/article/the-10-companies-hiring-more-software-engineers-than-anyone-else-in-silicon-valley/> and <https://www.monster.com/career-advice/article/top-tech-employers-job-listings>.

¹⁶We also verified that Payscale.com’s estimates were comparable to those on Glassdoor.com, a similar website collecting salary data. For example https://www.payscale.com/research/U.S./Job=Software_Engineer/Salary/3f79787f/Amazon.com-Inc-Seattle-WA and https://www.glassdoor.com/Salary/Amazon-Software-Engineer-Salaries-E6036_D_K07,24.htm. The distribution of base salaries reported to these types of websites is quite similar to those reported to the U.S. Census. For example, Glassdoor.com has benchmarked its salary data against Census data and published the results several times, and they are remarkably similar for base pay (Glassdoor, 2019).

¹⁷For IBM, which had no software engineer salary data in its headquarters of Armonk, N.Y. we instead used salaries from its other major campus in San Jose, California.

¹⁸The actual wage gap is difficult to compute, and is beyond the scope of this paper. Publicly available salary data about specific firms—including the sources we used above (Glassdoor.com and Payscale.com) and all others we consulted—do not contain gender-specific wage values.

¹⁹We restrict the ACS data to individuals with a bachelor’s degree (only) who are employed either in computer occupations (ACS Occupation Codes 10XX and 11XX) or specifically as computer software engineers (ACS Occupation Code 1020). Note that our Payscale.com data combines data for men and women. On average, in the ACS, men in both computer and specifically software engineer occupations make 1.05 times the overall average. For computer occupations, women make on average 0.81 times what men make; for software developer occupations, women make on average 0.78 times what men make.

adjustments: we added a small amount of noise²⁰ and rounded to the nearest \$1,000. The noise and rounding produced only trivial changes to the distribution of salaries. However, it guaranteed that the “roundness” of disclosed salary numbers was randomly assigned and uncorrelated with a candidate’s gender, current employer, or other characteristics. Prior research suggests that round numbers are received differently in negotiation (Mason et al., 2013).

Each applicant had one job after graduation before his or her current job, as well as a college internship. Two jobs since graduation are typical, considering our candidates were in the full-time workforce for four to five years by the time of their applications.²¹ We injected small amounts of random variation in the start date and duration of the first job. This was in order to create realistic variation across candidates so they did not all contain identical dates. The post college job started shortly after college graduation and had a total tenure of between 6 and 17 months (randomly selected). The duration of the current job varied by when the recruiter viewed the applicant’s materials, but all the current jobs started between February 2014 and November 2015.

The applications also listed additional skills, achievements and coursework. We modeled these details after the résumés of real software engineers based on discussions with real employers in this industry, and based on the profiles and resumes of candidates at the major employers and companies listed above (i.e., the fictional candidates’ real-life counterparts).

F.1 Candidate Characteristics Balance

We have full control of all the attributes of the job candidates, including whether they disclose, so we made sure to balance our candidates on attributes we were not primarily interested in. For example, the average year of graduation was 2013, and the proportion currently working at Amazon is 6% for candidates who don’t disclose as well as for those who do disclose, as shown in Table F1.

Table F1: Candidate Balance

	Female Candidate	Median Salary Cur- rent Empl (10K)	College Grad. Year	Disclosed Salary Cur- rent Empl (10K)	Amazon	Facebook	IBM	% of Sam- ple
All Candidates	0.50	9.97	2013.66	9.71	0.06	0.13	0.09	100.0
No Salary Disclosed	0.50	9.98	2013.67	.	0.06	0.13	0.10	40.6
Salary Disclosed	0.50	9.97	2013.65	9.71	0.06	0.12	0.09	59.4
No Salary Prompt	0.50	9.97	2013.66	9.71	0.06	0.13	0.09	43.8
Has Salary Prompt	0.50	9.97	2013.66	9.71	0.06	0.13	0.09	56.3

Notes: This table shows the attributes of the fictitious job candidates overall, by whether their application included a salary history prompt, and by whether the candidate disclosed their salary in the application form. These are balanced by design.

²⁰This draws from a uniform random distribution from -\$2,000 to +\$2,000 in \$1,000 increments.

²¹According to the BLS, median job tenure for those 20-24 is 1.3 years, and for those 25-34 is 2.8 years (<https://www.bls.gov/news.release/tenure.t01.htm>).

G Spillover Controls

In our analysis, we include spillover terms that take into account the disclosures of other candidates who were included in the same packet of eight sent to a recruiter. Each line of the data is a single job candidate. We control for the average of all the other salaries disclosed amongst the eight, excluding the job candidate's own, and fixed effects for the subtreatment the packet was assigned.²²

H Disclosure Indifference Salaries

We use the coefficients derived from Equation 4 to impute salaries where (on average) candidates would be indifferent between disclosing and not. In this setup, the candidate is indifferent at the point where $y_{ij}(1, h)$ (their outcome from disclosing a salary of h) is equal to $y_{ij}(0,)$ (their outcome from silence). Assuming a linear data generating process akin to Equation 4, then:

$$y_{ij}(1, h) = \beta_1 + \beta_2(h - \mu) + v_i + \gamma_i + \beta_3[SpilloverControls_{ij}] + \beta_4[RecruiterControls_j] + \epsilon_j \quad (6)$$

where μ is the average disclosed wage used to de-mean the disclosed salaries (so that $h - \mu$ is the demeaned disclosure amount). In our case, μ is the gender-specific mean of disclosed salary amounts within firm. The spillover terms represent how other candidates' disclosures affect the focal candidate. We assume these are not affected differently for disclosing and non-disclosing candidates. Similarly for $y_{ij}(0, \emptyset)$, we have that:

$$y_{ij}(0, \emptyset) = v_i + \gamma_i + \beta_3[SpilloverControls_{ij}] + \beta_4[RecruiterControls_j] + \epsilon_j \quad (7)$$

Setting these equal yields $0 = \beta_1 + \beta_2(s - \mu)$. Rearranging, this implies that the indifference salary is equal to $h = -\beta_1/\beta_2 + \mu$. In other words, the salary level h at which candidates would be (on average) indifferent about disclosing (or revealing s) is the one where the intercept effect of disclosing (β_1 in Equation 4) is fully counteracted by the slope effects coming from the amount ($\beta_2(h - \mu)$ from Equation 4). To find this value, we estimate Equation 4 on our data and calculate this value for h as a function of the estimated coefficients β_1 and β_2 (and the constant μ). For standard errors for this term, we use the delta method (Oehlert, 1992; Rice, 2006).

Gender-Specific Average Indifference Salaries. The above describes how we calculate a single average indifferent point. We also calculate these separately for men and women. This can either be performed in separate on male/female regressions. Alternatively, we can use

²²For example, if the first line of the data is Jessica and she discloses \$105,000, and 3 other people in the packet disclose \$90,000, \$97,000, and \$103,000 then we include the average of those three disclosures (\$97,000) as a control variable.

the augmented version of Equation 4 described in the main text that includes gender interactions (both with $Disclosed_{ij}$ and $Disclosed_{ij} \times AmountDisclosedDemeaned_{ij}$). Specifically, the interactions are with a $FemaleCandidate_i$ dummy variable, as shown below.

$$\begin{aligned}
y_{ij} = & \beta_1 Disclosed_{ij} + \beta_2 Disclosed_{ij} \times AmountDisclosedDemeaned_{ij} \\
& + \beta_3 FemaleCandidate_i \times Disclosed_{ij} + \\
& + \beta_4 FemaleCandidate_i \times Disclosed_{ij} \times AmountDisclosedDemeaned_{ij} \\
& + v_i + \gamma_j + \beta_3 [SpilloverControls_{ij}] + \beta_4 [RecruiterControls_j] + \epsilon_j
\end{aligned} \tag{8}$$

For the male version we calculate $h = -\beta_1/\beta_2 + \mu$. For the female version, we add the female-interaction coefficients for both the slope and interaction terms. As such, the indifference salary for women would be $h = -(\beta_1 + \beta_3)/(\beta_2 + \beta_4) + \mu$. Note that the μ term is different for men and women because we de-meaned by gender. As before, for standard errors we used the delta method (Oehlert, 1992; Rice, 2006). These values tell us the salaries below which the average candidate in either gender would be better off being silent, and above which the average candidate would be better off disclosing. We then compare these values to the distribution of female and male salaries in our experiment to understand the costs and benefits of disclosure.

I Recruiter Knowledge of Average Market Wages

One potential alternative explanation for our results on the effects of silence is that recruiters simply misjudged the average level of market wages for this job. Our subjects may have believed that silent workers earned market-average wages but misjudged average-market pay levels for software engineers. Our candidates' disclosure amounts were based on third-party data about true, accurate market levels, and our recruiter subjects were experienced professionals. Insofar as they were not, they could estimate market levels using the same publicly available tools. In fact, we administered a brief questionnaire to the recruiters after they completed the main task, and we found that when recruiters were presented with packets with no disclosed salaries they were more likely to report doing external research to help determine salary levels (82% versus 73.5% for those who saw zero rather than four or eight disclosed salaries, one-sided $p=0.09$).²³ This concords with the findings of Barach and Horton (2021), which shows that when employers could not observe full compensation histories, they asked applicants more questions and spent more time acquiring additional information.

Nonetheless, they may have underestimated market wages for software engineers. To address this, Tables K8, K9, and K10 examine the subset of recruiters who receive packets of half-disclosing, half-silent candidates. These subjects address this question because the half of candidates who disclosed a number gave a reminder of general market wages to use as

²³The question asked "How did you make judgments on the salary related questions? Select all that apply", and the options were "Used my previous experience with salaries in this setting"; "Looked up salaries on a website like payscale.com, glassdoor.com, etc."; "Spoke with others familiar with salaries for software engineers"; "Other". We considered the recruiter to "do research" if he or she reported looking up salaries or speaking with others.

a benchmark for the silent candidates. However, in this sample, our results are very similar to the full sample—silent candidates are assumed to be adversely selected. This suggests our result is not likely an artifact of recruiter inexperience or lack of knowledge of market wages.

J External Validity: SANS conditions

This section discusses our findings using the SANS conditions (selection, attrition, naturalness, and scalability) suggested by [List \(2020\)](#).

Selection. Our subject pool of recruiters is broadly representative of the target population of recruiters, including those in software (Appendix C). Our candidates and job openings come from the market for software engineers. The candidates were based on the actual job applications for these positions. We choose to study software engineering in part because of the persistent gender disparities in this industry (Appendix A). The task assigned to recruiters – to suggest both candidates and wages – is performed by all employers. Section 6.1 reviews the specific practices used by businesses for this task. We used a set of hiring materials – from the job description and application to the recruiter instructions – based on those at real companies in this industry.

Attrition. About 17% of subjects did not complete the task after being sent materials. Dropout from the study was not correlated with our randomly assigned treatment arms (Table J1). Additionally, we include tables that explore if attrition is predicted by observable characteristics of the recruiter or the job candidates in Tables J2 and J3.

Table J1: Attrition

	Final Data Set	Final Data Set	Final Data Set
Salary Prompt Appears On Job Application Form	.019 (.043)		.026 (.044)
At Least One Disclosure		-.069 (.047)	-.073 (.048)
Observations	307	307	307
R ²	.00061	.0056	.0069
Mean Dep. Var	.83		.83

Notes: This table studies attrition of invited subjects into the study. Of the 307 subjects sent materials, 256 gave a full set of valid answers and ended up in the final dataset. In the regressions above, we study which invited subjects remained in the study as a function of their treatment assignment.

We find no evidence of differential attrition by observable characteristics of the recruiter (Column 1 in Table J2). In addition, we find no patterns of differential attrition when we add in the major treatment statuses (Column 2 in Table J2). In Column 3 of Table J2, we study the interactions. Although the standard errors on the interacted coefficients are wider, we similarly find no pattern of differential attrition among subjects sent the study materials. In Table J3, we also tested for differential attrition based on the characteristics of the candidates who were included in the packet of job applications. We similarly see no patterns of differential attrition based on candidate characteristics, or their interaction with the prompt or whether a salary was disclosed.

Table J2: Attrition By Recruiter Observable Characteristics

	FinalDataSet	FinalDataSet	FinalDataSet
Female Recruiter (Guess)	-0.01 (0.05)	-0.01 (0.05)	-0.18 (0.13)
White (Guess)	0.06 (0.07)	0.07 (0.07)	0.20 (0.18)
Black (Guess)	0.07 (0.08)	0.07 (0.08)	0.08 (0.21)
Hourly Rate	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Avg Feedback Score	0.01 (0.01)	0.01 (0.01)	0.02 (0.04)
Total Feedback Scores	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Total Passed Tests	-0.02 (0.02)	-0.02 (0.02)	-0.07 (0.08)
Total Portfolio Items	0.02 (0.02)	0.02 (0.02)	0.04 (0.07)
Total Hours Billed	-0.00 (0.00)	-0.00 (0.00)	0.00 (0.00)
Salary Prompt Appears On Job Application Form		0.03 (0.04)	0.05 (0.18)
At Least One Disclosure		-0.07 (0.05)	-0.19 (0.21)
Female Recruiter (Guess) * HasAPrompt			0.08 (0.11)
White (Guess) * HasAPrompt			-0.21 (0.16)
Black (Guess) * HasAPrompt			-0.25 (0.17)
Hourly Rate * HasAPrompt			0.00 (0.00)
Avg Feedback Score * HasAPrompt			0.02 (0.03)
Total Feedback Scores * HasAPrompt			-0.01 (0.01)
Total Passed Tests * HasAPrompt			0.02 (0.04)
Total Portfolio Items * HasAPrompt			-0.02 (0.05)
Total Hours Billed * HasAPrompt			-0.00 (0.00)
Female Recruiter (Guess) * AnyDisc			0.16 (0.14)
White (Guess) * AnyDisc			0.01 (0.19)
Black (Guess) * AnyDisc			0.20 (0.22)
Hourly Rate * AnyDisc			-0.00 (0.00)
Avg Feedback Score * AnyDisc			-0.03 (0.05)
Total Feedback Scores * AnyDisc			0.00 (0.00)
Total Passed Tests * AnyDisc			0.05 (0.08)
Total Portfolio Items * AnyDisc			0.00 (0.08)
Total Hours Billed * AnyDisc			-0.00 (0.00)
Observations	307.00		307.00
R ²	0.02		0.07
Mean Dep. Var	0.83		0.83

Notes: This table studies the attrition of invited subjects into the study. Of the 307 subjects sent materials, 256 were in the final data set. In the regressions above, we study whether observable attributes of invited subjects interacted with whether they were assigned to a packet which HasAPrompt for salary history information or whether any of the candidates disclosed a salary (AnyDisc) are predictive of if they ended up in the final data set. For the variables Female (Guess), White (Guess), and Black (Guess) this is our best guess of a recruiter's gender and race based on name and photo. Hourly rate is the hourly rate the recruiter charges for their services. Avg Feedback Score, Total Feedback Scores, Total Passed Tests, Total Portfolio Items, and Total Hours Billed are posted on the freelancing platform and measured before we hire the recruiter. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table J3: Attrition By Job Candidate Characteristics

	Final Data Set	Final Data Set	Final Data Set
Female Candidate	.000068 (.0018)	.00038 (.0018)	.002 (.0074)
Median Salary Current Empl (10K)	.00055 (.0019)	.00054 (.0019)	-.0036 (.006)
Grad Year	.011 (.01)	.0099 (.01)	.0002 (.018)
Amazon	.02 (.023)	.02 (.023)	.018 (.039)
Facebook	.0077 (.0065)	.0067 (.0065)	.012 (.024)
IBM	.028 (.018)	.028 (.018)	-.0082 (.029)
HasAPrompt		.024 (.043)	-.39 (44)
SalaryWasDisclosed		-.034 (.032)	2.7 (55)
FemaleCandidate * HasPrompt			-.00016 (.0046)
PrevEmployerMedianSalary * HasPrompt			-.00023 (.004)
SchoolGradYear * HasPrompt			.019 (.022)
Amazon * HasPrompt			.04 (.048)
Facebook * HasPrompt			.02 (.016)
IBM * HasPrompt			.051 (.037)
FemaleCandidate * Disclosed			-.004 (.015)
PrevEmployerMedianSalary * Disclosed			.0071 (.0094)
SchoolGradYear * Disclosed			-.0014 (.027)
Amazon * Disclosed			-.032 (.052)
Facebook * Disclosed			-.028 (.043)
IBM * Disclosed			.015 (.048)
Observations	2,456	2,456	2,456
R ²	.00083	.0035	.0045
Mean Dep. Var	.83	.83	.83

Notes: This table studies the attrition of invited subjects into the study depending on the job candidates they were shown. Of the 307 subjects sent materials, 256 gave a full set of valid answers and ended up being in the final data set. In the regressions above, we study whether attributes of fictitious job candidates interacted with whether they filled out an application form which HasAPrompt for salary history information or whether the candidates disclosed their salary (Disclosed) are predictive of if they ended up in the final data set. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Naturalness. Our field experiment engages recruiters in an organic setting for their jobs. This may be important because subjects in laboratory experiments may be tempted to behave more benevolently than they would in reality, particularly if they sense their discrimination is being measured.

We asked the salary history question on the job application. According to our survey, this is the most common way to ask. However, we acknowledge that salary questions could also arise interactively during an interview. In this context, candidates and employers could exchange additional information to clarify the interpretation of the salary history. In a setting like ours, a candidate whose previous salary is “too high” could clarify their expectations and potentially avoid rejection.

We do not capture these effects. Because our experiment does not feature these clarifying questions, we isolate the effect of the salary history information and separate it from the disclosure of additional information (such as expectations or other mitigating circumstances). Communication of salary expectations is a separate, rich topic (Roussille, [forthcoming](#)). Some of our results suggest reasons why such clarifications would be useful: Absent clarification, salary histories alone could be either “too high” or “too low.” These types of additional, clarifying disclosures would provide important mitigating effects for any policy implications (at least within interview settings).

Scalability. Although our study was motivated by a public policy question, our primary aim is to measure one of the main ingredients to the policy effects: What employers learn from disclosures (and the lack thereof). Some aspects of salary history bans can be scaled. Bans appear to be effective at reducing or eliminating certain questions.

However, our analysis suggests that other aspects are clearly not scalable. Although ban legislation can stop all employers from asking, they cannot stop all candidates from volunteering. They also cannot stop employers from guessing why certain candidates are not volunteering. Our conceptual model shows how unraveling would proceed, and our empirical results contain some evidence of unraveling dynamics (i.e., silence assumed to be a negative signal).

Outside of our experiment, we find additional suggestive evidence of unraveling. A report by the New York Times about salary history bans in 2021 says, “some people [...] volunteered their salary history.” One candidate told the Times. “I prefer to be direct about what I’m making.”²⁴ In our surveys of the American workforce on this topic (Cowgill et al., 2022), we found a 10 percentage point increase in the number of people volunteering their salary unprompted between two waves of our survey (November 2019 and May 2021). Unraveling and other adjustments to bans raise the potential for a “voltage drop” after scaling up this policy.

²⁴<https://www.nytimes.com/2021/12/30/business/salary-negotiation-pay.html>

K Additional Empirical Analysis

K.1 Additional Outcomes

Table K1: Additional Outcomes Prompted Salary Disclosures (Similar to Table 2)

	(1) Outside Option 5th %tile	(2) Outside Option 5th %tile	(3) Outside Option 95th %tile	(4) Outside Option 95th %tile	(5) ≥ 2 Other Offers	(6) ≥ 2 Other Offers
Disclosed	0.70*** (0.12)	0.73*** (0.19)	0.26 (0.34)	0.72*** (0.19)	0.06 (0.05)	0.15* (0.07)
Disc x Amount (Demeaned)	0.67*** (0.05)	0.67*** (0.08)	0.72*** (0.12)	0.67*** (0.15)	0.01 (0.02)	0.01 (0.02)
Disc x Prompt		-0.05 (0.24)		-0.93 (0.67)		-0.16+ (0.09)
Disc x Amount (DM) x Prompt		-0.01 (0.09)		0.09 (0.10)		0.01 (0.03)
Mean Non-Disclosers (Prompt):	8.16 8.23 8.09		11.01 11.43 10.59		0.54 0.57 0.50	
Non-Disclosers (No Prompt):		8.07		10.65 10.77 10.53		0.51 0.54 0.47
²	0.53		0.12	0.12	0.05	0.05
Observations	2048		2048	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2 and Appendix E. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table K2: Additional Outcomes Gender, Disclosure and Amount (Similar to Table 3)

	(1) Outside Option 5th %tile	(2) Outside Option 5th %tile	(3) Outside Option 95th %tile	(4) Outside Option 95th %tile	(5) ≥ 2 Other Offers	(6) ≥ 2 Other Offers
Disclosed	1.17*** (0.15)	1.25*** (0.22)	0.39 (0.52)	1.14*** (0.24)	0.06 (0.05)	0.15* (0.07)
Female x Disclosed	-0.95*** (0.13)	-1.05*** (0.18)	-0.27 (0.41)	-0.84*** (0.23)	0.04 (0.05)	-0.03 (0.07)
Disc x Amount (Demeaned)	0.71*** (0.06)	0.75*** (0.09)	0.74*** (0.19)	0.73*** (0.22)	0.01 (0.02)	0.01 (0.02)
F x Disc x Amount (DM)	0.03 (0.04)	-0.05 (0.05)	-0.01 (0.15)	-0.10 (0.16)	0.00 (0.02)	0.01 (0.03)
Disc x Prompt		-0.15 (0.29)		-1.54 (1.00)		-0.16+ (0.09)
F x Disc x Prompt		0.21 (0.24)		1.23 (0.83)		0.11 (0.09)
F x Prompt		-0.16 (0.16)		-1.29+ (0.78)		-0.01 (0.06)
Disc x Amount (DM) x Prompt		-0.06 (0.11)		0.02 (0.12)		0.01 (0.03)
F x Disc x Amount (DM) x Prompt		0.12* (0.06)		0.15+ (0.08)		-0.01 (0.03)
Mean Non-Disclosers (Prompt):						
<i>All</i>	8.16		11.01		0.54	
<i>Male</i>	8.23		11.43		0.57	
<i>Female</i>	8.09		10.59		0.50	
Mean Non-Disclosers (No Prompt):						
<i>All</i>		8.07		10.65		0.51
<i>Male</i>		8.11		10.77		0.54
<i>Female</i>		8.04		10.53		0.47
R ²	0.54	0.54	0.12	0.12	0.05	0.05
Observations	2048	2048	2048	2048	2048	2048

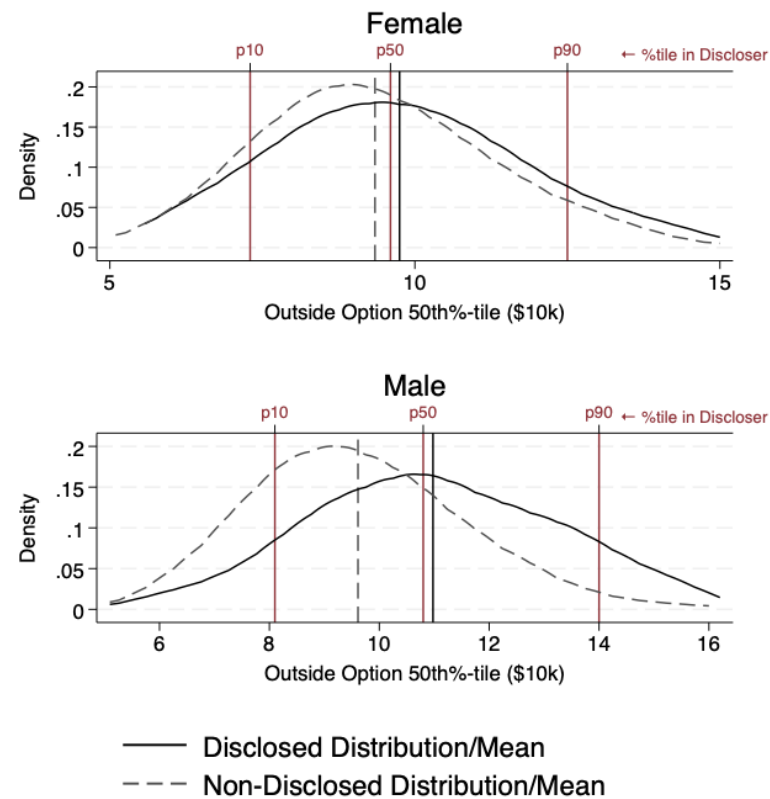
Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2 and Appendix E. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K3: Additional Outcomes Average Effect of Disclosing a High versus Low Salary (Similar to Table 4)

	(1) Outside Option 5th %tile	(2) Outside Option 95th %tile	(3) ≥ 2 Other Offers
Disclosed 25th %ile Salary	0.40** (0.13)	-0.47 (0.50)	0.05 (0.05)
Female \times Disclosed 25%tile Salary	-0.58*** (0.11)	0.16 (0.39)	0.04 (0.05)
Disclosed 75th %ile Salary	1.53*** (0.16)	0.83+ (0.50)	0.06 (0.05)
Female \times Disclosed 75%tile Salary	-0.50*** (0.12)	0.11 (0.40)	0.05 (0.05)
Female 25th Disclosure Effect:			
<i>Total</i>	-0.18	-0.31	0.10
<i>p-value</i>	0.20	0.32	0.06
Female 75th Disclosure Effect:			
<i>Total</i>	1.03	0.94	0.11
<i>p-value</i>	0.00	0.00	0.02
Mean Male Non-Disclosers	8.23	11.43	0.57
Mean Female Non-Disclosers	8.09	10.59	0.50
R ²	0.48	0.10	0.05
Observations	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2 and Appendix E. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Figure K1: Distribution of Outcomes by Disclosure Status (Similar to Figure 2)



Notes: This figure shows the distribution of recruiter's choices. The solid and dashed density plot represents choices in response to a given salary disclosure for male and female candidates (among the job applications containing a disclosure). While the solid and dashed vertical lines represent mean outcomes for disclosers and non-disclosers respectively. The light gray vertical lines mark the percentile of the outcome in the disclosers' distribution as listed.

K.2 Robustness Checks

Table K4: Conditional On Callback (Similar to Table 3)

	(1)	(2)	(3)	(4)	(5)	(6)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer
Disclosed	0.78*** (0.17)	1.05*** (0.24)	0.95*** (0.16)	1.00*** (0.22)	0.85*** (0.15)	1.03*** (0.23)
Female x Disclosed	-0.82*** (0.17)	-0.84*** (0.24)	-0.97*** (0.16)	-0.83*** (0.22)	-0.91*** (0.16)	-0.90*** (0.21)
Disc x Amount (Demeaned)	0.65*** (0.07)	0.70*** (0.10)	0.77*** (0.06)	0.79*** (0.09)	0.73*** (0.07)	0.78*** (0.09)
F x Disc x Amount (DM)	0.01 (0.06)	-0.04 (0.07)	0.00 (0.05)	-0.10 (0.07)	-0.03 (0.05)	-0.04 (0.07)
Disc x Prompt		-0.53 (0.33)		-0.12 (0.31)		-0.35 (0.30)
F x Disc x Prompt		0.02 (0.34)		-0.26 (0.30)		-0.06 (0.31)
F x Prompt		-0.09 (0.23)		0.05 (0.22)		0.08 (0.22)
Disc x Amount (DM) x Prompt		-0.07 (0.12)		-0.03 (0.10)		-0.08 (0.11)
F x Disc x Amount (DM) x Prompt		0.08 (0.09)		0.16* (0.08)		0.02 (0.08)
Mean Non-Disclosers (Prompt):						
<i>All</i>	10.87		9.79		10.17	
<i>Male</i>	10.95		9.87		10.23	
<i>Female</i>	10.79		9.70		10.12	
Mean Non-Disclosers (No Prompt):						
<i>All</i>		10.82		9.64		10.09
<i>Male</i>		10.86		9.71		10.15
<i>Female</i>		10.78		9.57		10.04
R ²	0.38	0.38	0.52	0.52	0.44	0.44
Observations	1297	1297	1297	1297	1297	1297

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K5: To Candidate If Only Make One Offer (Similar to Table 3)

	(1)	(2)	(3)	(4)	(5)	(6)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer
Disclosed	-0.06 (0.52)	0.58 (0.76)	0.24 (0.51)	1.02 (0.78)	0.19 (0.46)	0.63 (0.71)
Female x Disclosed	-0.14 (0.57)	-0.06 (0.80)	-0.41 (0.56)	-0.50 (0.70)	-0.38 (0.53)	-0.09 (0.73)
Disc x Amount (Demeaned)	0.42+ (0.22)	0.71** (0.26)	0.57** (0.21)	0.80** (0.30)	0.50* (0.21)	0.85*** (0.23)
F x Disc x Amount (DM)	0.07 (0.23)	-0.05 (0.28)	0.20 (0.22)	0.03 (0.30)	0.12 (0.21)	-0.15 (0.27)
Disc x Prompt		-1.48 (0.98)		-1.63 (0.98)		-1.02 (0.90)
F x Disc x Prompt		-0.11 (1.00)		0.11 (0.94)		-0.57 (0.94)
F x Prompt		0.27 (0.85)		0.46 (0.83)		0.63 (0.79)
Disc x Amount (DM) x Prompt		-0.47 (0.29)		-0.38 (0.32)		-0.55* (0.27)
F x Disc x Amount (DM) x Prompt		0.07 (0.28)		0.14 (0.29)		0.27 (0.27)
Mean Non-Disclosers (Prompt):						
<i>All</i>	11.88		10.39		10.90	
<i>Male</i>	12.09		10.54		11.04	
<i>Female</i>	11.59		10.20		10.70	
Mean Non-Disclosers (No Prompt):						
<i>All</i>		11.80		10.23		10.84
<i>Male</i>		12.11		10.45		11.18
<i>Female</i>		11.44		9.98		10.46
R ²	0.32	0.33	0.40	0.41	0.36	0.37
Observations	256		256	256	256	256

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K6: Recruiter Doesn't Think Candidates Misreport (Similar to Table 3)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer	Call- back	Call- back
Disclosed	0.73*** (0.19)	1.13*** (0.23)	0.97*** (0.19)	1.21*** (0.25)	0.83*** (0.18)	1.18*** (0.22)	-0.00 (0.05)	-0.01 (0.07)
Female x Disclosed	-0.40* (0.18)	-0.78*** (0.23)	-0.76*** (0.16)	-0.99*** (0.21)	-0.49** (0.18)	-0.92*** (0.21)	0.15** (0.05)	0.12+ (0.07)
Disc x Amount (Demeaned)	0.58*** (0.08)	0.64*** (0.11)	0.69*** (0.08)	0.70*** (0.10)	0.61*** (0.08)	0.69*** (0.10)	-0.05* (0.02)	-0.07* (0.03)
F x Disc x Amount (DM)	0.03 (0.06)	0.01 (0.08)	0.01 (0.05)	-0.03 (0.07)	0.05 (0.05)	0.06 (0.07)	0.04* (0.02)	0.06* (0.03)
Disc x Prompt		-0.86* (0.38)		-0.57 (0.37)		-0.72* (0.35)		0.00 (0.11)
F x Disc x Prompt		0.80* (0.34)		0.54+ (0.30)		0.87** (0.33)		0.07 (0.11)
F x Prompt		-0.52* (0.25)		-0.57* (0.23)		-0.44+ (0.25)		-0.07 (0.08)
Disc x Amount (DM) x Prompt		-0.12 (0.13)		-0.01 (0.12)		-0.17 (0.12)		0.03 (0.04)
F x Disc x Amount (DM) x Prompt		0.04 (0.09)		0.08 (0.08)		-0.01 (0.08)		-0.03 (0.03)
Mean Non-Disclosers (Prompt):								
All	10.34		9.55		9.70		0.63	
Male	10.53		9.68		9.85		0.66	
Female	10.15		9.42		9.54		0.61	
Mean Non-Disclosers (No Prompt):								
All		10.33		9.38		9.66		0.62
Male		10.43		9.39		9.74		0.63
Female		10.23		9.37		9.59		0.61
R ²	0.39	0.39	0.53	0.53	0.44	0.44	0.05	0.05
Observations	1264	1264	1264	1264	1264	1264	1264	1264

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K7: Average Effect of Disclosing by Salary Amount and if Recruiters in Ban States

	(1)	(2)	(3)	(4)	(5)	(6)
	WTP	WTP	Outside Option 50th %tile	Outside Option 50th %tile	Offer	Offer
Disclosed	0.52*** (0.13)	0.50*** (0.14)	0.61*** (0.13)	0.61*** (0.13)	0.58*** (0.12)	0.54*** (0.13)
Disc x Amount (Demeaned)	0.64*** (0.05)	0.61*** (0.06)	0.73*** (0.05)	0.73*** (0.05)	0.66*** (0.05)	0.64*** (0.05)
RecruiterInBan x Disclosed		0.07 (0.15)		0.03 (0.14)		0.12 (0.13)
RecruiterInBan x Amount Demeaned		0.09 (0.06)		-0.00 (0.05)		0.08 (0.05)
R ²	0.38	0.38	0.52	0.52	0.43	0.43
Observations	2048	2048	2048	2048	2048	2048

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4 and similar to those shown in Table 2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

K.3 Half Disclosing

Table K8: Half Disclosed (Similar to Table 2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer	Call- back	Call- back
Disclosed	0.54*** (0.13)	0.75*** (0.19)	0.64*** (0.12)	0.75*** (0.18)	0.60*** (0.11)	0.74*** (0.18)	-0.05 (0.05)	0.00 (0.07)
Disc x Amount (Demeaned)	0.63*** (0.07)	0.62*** (0.11)	0.64*** (0.07)	0.63*** (0.10)	0.65*** (0.07)	0.67*** (0.10)	-0.04 (0.03)	-0.03 (0.03)
Disc x Prompt		-0.43 (0.26)		-0.22 (0.24)		-0.28 (0.23)		-0.09 (0.11)
Disc x Amount (DM) x Prompt		0.02 (0.14)		0.02 (0.13)		-0.04 (0.13)		-0.01 (0.04)
Mean Non-Disclosers (Prompt):								
<i>All</i>	10.34		9.64		9.72		0.63	
<i>Male</i>	10.53		9.72		9.87		0.66	
<i>Female</i>	10.15		9.56		9.57		0.59	
Mean Non-Disclosers (No Prompt):								
<i>All</i>		10.15		9.41		9.55		0.64
<i>Male</i>		10.24		9.44		9.65		0.64
<i>Female</i>		10.07		9.37		9.45		0.64
R ²	0.27	0.27	0.35	0.35	0.31	0.31	0.06	0.06
Observations	768	768	768	768	768	768	768	768

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K9: Half Disclosed (Similar to Table 3)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Option	Outside Option	Offer	Offer	Call- back	Call- back
Disclosed	0.70*** (0.20)	1.01*** (0.27)	1.04*** (0.17)	1.16*** (0.25)	0.83*** (0.17)	1.07*** (0.23)	-0.05 (0.05)	0.00 (0.07)
Female x Disclosed	-0.32 (0.25)	-0.53 (0.33)	-0.81*** (0.20)	-0.83** (0.31)	-0.46* (0.22)	-0.65* (0.28)	0.15* (0.06)	0.10 (0.09)
Disc x Amount (Demeaned)	0.61*** (0.09)	0.60*** (0.14)	0.65*** (0.08)	0.69*** (0.12)	0.62*** (0.08)	0.63*** (0.12)	-0.04 (0.03)	-0.03 (0.03)
F x Disc x Amount (DM)	0.06 (0.09)	0.08 (0.11)	0.04 (0.08)	-0.06 (0.11)	0.08 (0.08)	0.10 (0.10)	0.05+ (0.02)	0.06+ (0.03)
Disc x Prompt		-0.63 (0.39)		-0.24 (0.33)		-0.48 (0.34)		-0.09 (0.11)
F x Disc x Prompt		0.41 (0.49)		0.04 (0.39)		0.39 (0.44)		0.11 (0.13)
F x Prompt		-0.43 (0.33)		-0.17 (0.26)		-0.20 (0.28)		-0.14 (0.09)
Disc x Amount (DM) x Prompt		0.03 (0.17)		-0.07 (0.16)		-0.03 (0.16)		-0.01 (0.04)
F x Disc x Amount (DM) x Prompt		-0.03 (0.14)		0.19 (0.13)		-0.04 (0.13)		-0.02 (0.04)
Mean Non-Disclosers (Prompt):								
<i>All</i>	10.34		9.64		9.72		0.63	
<i>Male</i>	10.53		9.72		9.87		0.66	
<i>Female</i>	10.15		9.56		9.57		0.59	
Mean Non-Disclosers (No Prompt):								
<i>All</i>		10.15		9.41		9.55		0.64
<i>Male</i>		10.24		9.44		9.65		0.64
<i>Female</i>		10.07		9.37		9.45		0.64
R ²	0.27	0.27	0.36	0.35	0.32	0.31	0.06	0.06
Observations	768	768	768	768	768	768	768	768

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

Table K10: Half Disclosed (Similar to Table 4)

	(1)	(2)	(3)	(4)
	WTP	Outside Option	Offer	Callback
Disclosed 25th %ile Salary	-0.02 (0.20)	0.20 (0.17)	0.08 (0.17)	-0.01 (0.07)
Female × Disclosed 25%tile Salary	-0.14 (0.25)	-0.51* (0.20)	-0.27 (0.21)	0.09 (0.08)
Disclosed 75th %ile Salary	1.13*** (0.26)	1.59*** (0.21)	1.28*** (0.22)	-0.07 (0.07)
Female × Disclosed 75%tile Salary	-0.01 (0.30)	-0.59* (0.25)	-0.15 (0.27)	0.20* (0.08)
Female 25th Disclosure Effect:				
<i>Total</i>	-0.17	-0.31	-0.19	0.08
<i>p-value</i>	0.37	0.09	0.27	0.20
Female 75th Disclosure Effect:				
<i>Total</i>	1.12	1.00	1.14	0.13
<i>p-value</i>	0.00	0.00	0.00	0.04
Mean Male Non-Disclosers	10.53	9.72	9.87	0.66
Mean Female Non-Disclosers	10.15	9.56	9.57	0.59
R ²	0.21	0.30	0.25	0.06
Observations	768	768	768	768

Notes: All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in Section 3.2. Outcomes measured in dollars (e.g. Outside Option) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

L Experimental Materials

L.1 Sample Job Application: Salary History Asked + Candidate Discloses

Samantha [REDACTED]

Application Details for Software Engineering position

All sections are **required** except where noted. For candidates who are interviewed, all information entered below will be verified.

Candidate Information

Candidate Id: 774

Mailing Address: [REDACTED]

City/State: [REDACTED]

ZIP: [REDACTED]

Phone: ([REDACTED]) [REDACTED]-[REDACTED]

Email: [REDACTED]

URL: http://[REDACTED]

Are you legally authorized to work in the US? Y

Are you over the age of 18?: Y

Are you willing to relocate for this position? Y

Will you now (or in the future) require visa sponsorship? N

Employment History (Last Three Jobs)

Title: Software Engineer

Company Name: IBM

Location: San Jose, CA

Dates: 01/2015 - Present

Position Description, Duties, Responsibilities:

* Developing and implementing new feedback system for user concerns, bugs, and defect tracking regarding use and functionality of new interfaces.

* Coding web designed interfaces using Java, XML, XSL, AJAX, and JWS.

* Implement the command-line interface for the Universal Authentication Protocol (UAP) in E-directory.

Title: Software Developer

Company Name: Amazon

Location: Seattle, WA

Dates: 05/2014 - 01/2015

Position Description, Duties, Responsibilities:

* Developed code and unit tests in Python for server-side and in JavaScript for web components.

* Deployed and tested code on Linux-based EC2 instances in a distributed AWS cloud environment.

* Created and maintained automated jobs to build and test software.

* Developed and implemented working plans for the formulation of front and back-end web applications.

* Developed various algorithms to mitigate program interference.

Title: Programming Intern

Company Name: Intraix

Location: Ayer Rajah Crescent, SG

Dates: 05/2013 - 08/2013

Position Description, Duties, Responsibilities:

Automated black box and white box tests for an Android application "Klug," using Appium and Espresso framework. This helped developers expand features without much worry of breaking current functionalities.

Salary History

Annual Base Salary at Current or Most Recent Job: \$96,000

Education

Institution: Georgia Institute of Technology

Location: Atlanta, GA

Dates: 2010 - 2014

Graduated? Y

Level: BS (Bachelor of Science)

Subject/Major: Computer Science

Relevant Coursework:

Database and Information Management Systems, Java, Analysis of Algorithms, Data Systems, Matlab for Programmers, and Compiler Design

Additional Skills and Information

Experience developing in Java, HTML/CSS, JavaScript, Node.js, Ruby, Ruby on Rails, Shell, Python, SQL, LATEX.

L.2 Sample Job Application: Salary History Asked + Candidate Does Not Disclose

Christopher [REDACTED]

Application Details for Software Engineering position

All sections are **required** except where noted. For candidates who are interviewed, all information entered below will be verified.

Candidate Information

Candidate Id: 721

Mailing Address: [REDACTED]

City/State: [REDACTED]

ZIP: [REDACTED]

Phone: ([REDACTED]) [REDACTED]

Email: [REDACTED]

URL: http://[REDACTED]

Are you legally authorized to work in the US? Y

Are you over the age of 18? Y

Are you willing to relocate for this position? Y

Will you now (or in the future) require visa sponsorship? N

Employment History (Last Three Jobs)

Title: Programmer

Company Name: Apple

Location: Cupertino, CA

Dates: 10/2015 - Present

Position Description, Duties, Responsibilities:

Research, design, and implement scalable applications for information identification, extraction, analysis, retrieval, and indexing. Direct software design and development while remaining focused on client needs. Collaborate closely with other team members to plan, design, and develop robust solutions. Maintain front-end admin interface as well as back data processing.

Title: Programmer

Company Name: Verizon Communications, Inc.

Location: New York, NY

Dates: 07/2014 - 10/2015

Position Description, Duties, Responsibilities:

Designed, developed, and integrated software with test systems hardware for test engineering applications. Supported the design and testing of space systems software in all program phases, from initial design through coding, testing, and integration. Member of team responsible for developing a new high-end software package. Led team of 3 engineers to manage Windows client (C++) including feature development, debugging, and update release.

Title: Summer Programming Associate

Company Name: Facebook

Location: Menlo Park, CA

Dates: 06/2013 - 08/2013

Position Description, Duties, Responsibilities:

Intern on the Sales Platform team within Core Ads, which deals primarily with making tools to help salespeople make sales, usually by connecting them to advertisers. Worked on improving the infrastructure and data quality of our platform that helps sales teams find their clients. Languages/technologies: Hack (PHP), Python, Dataswarm.

Salary History (optional)

Annual Base Salary at Current or Most Recent Job:

Education

Institution: California Institute of Technology

Location: Pasadena, CA

Dates: 2010 - 2014

Graduated? Y

Level: BS (Bachelor of Science)

Subject/Major: Computer Science

Relevant Coursework:

Artificial language, hardware systems, analysis of algorithms. programming abstractions, data structures and algorithms

Additional Skills and Information

Production code launched using C/C++, Java, Javascript, Python, Perl. Back-end and research experience using Linux shell scripting, R, PiCloud/Multivac, Sawzall, MapReduce.

L.3 Sample Job Application: Salary History Not Asked + Candidate Does Not Disclose

Sarah [REDACTED]

Application Details for Software Engineering position

All sections are **required** except where noted. For candidates who are interviewed, all information entered below will be verified.

Candidate Information

Candidate Id: 1724 Mailing Address: [REDACTED] City/State: [REDACTED]
ZIP: [REDACTED] Phone: ([REDACTED]) [REDACTED]-[REDACTED] Email: [REDACTED] URL: http://[REDACTED]
Are you legally authorized to work in the US? Y Are you over the age of 18?: Y
Are you willing to relocate for this position? Y Will you now (or in the future) require visa sponsorship? N

Employment History (Last Three Jobs)

Title: Coder Company Name: Facebook Location: Menlo Park, CA Dates: 06/2014 - Present

Position Description, Duties, Responsibilities:

Enhancing existing web applications to meet current standards. Constructing complex queries using SQL in the IBM DB2 Database. Designing technical structure and modules for a new and better UX. Collaborating with senior developers to execute client work. Introducing automated acceptance and unit tests, while increasing coverage.

Title: Software Architect Company Name: Dell Location: Round Rock, TX Dates: 06/2013 - 06/2014

Position Description, Duties, Responsibilities:

Participate in application modification and development of new applications to meet business needs. Provide full life-cycle project expertise. Project work focused on business applications and e-business solutions. Responsibilities included application integration and development using .NET including C#, ASP.Net, WinForms, MS Exchange, and Microsoft Sharepoint Portal Server.

Title: Summer Coding Fellowship Company Name: Apple Location: Cupertino, CA Dates: 05/2012 - 08/2012

Position Description, Duties, Responsibilities:

Built an automated framework on the Apple Maps Team for validating the internal pipeline that manages how different layers of maps data integrate using Python.

Education

Institution: Cornell University Location: Ithaca, NY Dates: 2009 - 2013 Graduated? Y
Level: BS (Bachelor of Science) Subject/Major: Computer Science

Relevant Coursework:

Systems Programming and Machine Organization, Privacy and Technology, Data Science I, Networks, Computing Hardware, Cloud Computing.

Additional Skills and Information

Skills: JS, Java, XPages, Flex / AIR, Processing, Git, Eclipse, HTML.

L.4 Sample Job Application: Salary History Not Asked + Candidate Volunteers

Tyler [REDACTED]

Application Details for Software Engineering position

All sections are **required** except where noted. For candidates who are interviewed, all information entered below will be verified.

Candidate Information

Candidate Id: 621

Mailing Address: [REDACTED]

City/State: [REDACTED]

ZIP: [REDACTED]

Phone: ([REDACTED]) [REDACTED]

Email: [REDACTED]

URL: http://[REDACTED]

Are you legally authorized to work in the US? Y

Are you over the age of 18? Y

Are you willing to relocate for this position? Y

Will you now (or in the future) require visa sponsorship? N

Employment History (Last Three Jobs)

Title: Developer

Company Name: Amazon

Location: Seattle, WA

Dates: 02/2014 - Present

Position Description, Duties, Responsibilities:

- Develop automated REST API test cases to ensure proper error handling.
- Conduct regression tests on internal and external products and services in order to successfully integrate new solutions to existing systems.
- Review and approve code releases from development and marketing departments. ensure thorough client policy compliance.

Title: Coder

Company Name: Google

Location: Mountain View, CA

Dates: 05/2013 - 02/2014

Position Description, Duties, Responsibilities:

- Researched emerging technologies for database and network storage solutions by reviewing case studies and functionality to determine low-cost, but effective, models for supported environments.
- Provided leadership and decision making to impact infrastructure changes that included upgrading the Oracle database schema, applying new versions of Dart Enterprise, and implementing a virtualized hardware environment to reduce footprint and minimize data center presence.

Title: Software Development Trainee Company Name: GE Healthcare Location: Little Chalfont, UK Dates: 05/2012 - 08/2012

Position Description, Duties, Responsibilities:

Reduced waiting time to pull information from multiple systems - requests that used to take days, now only take minutes. Also worked closely with other IT professionals to design, test, and implement APIs in support of major ERP systems.

Education

Institution: University of Illinois at Urbana-Champaign

Location: Champaign, IL

Dates: 2009 - 2013

Graduated? Y

Level: BS (Bachelor of Science)

Subject/Major: Computer Science

Relevant Coursework:

C++, Java, Microprocessor systems, Cryptography, Human-computer interface technology, Computer networks, and Large scale systems

Additional Skills and Information

Skilled in Python (Django), Java, Ruby on Rails, JavaScript (AngularJS, jQuery), SQL, PHP, HTML, CSS. I make about \$125,000 per year right now (pre-bonus).

See <https://laurakgee.weebly.com/uploads/2/2/4/8/22485690/experimentalmaterialsforacg.pdf> for recruiter instructions and recruiter online evaluation form.

References

- Acemoglu, Daron and David Autor**, “Skills, tasks and technologies: Implications for employment and earnings,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 1043–1171.
- Acquisti, Alessandro and Christina M Fong**, “An experiment in hiring discrimination via online social networks,” 2015.
- Adler, Laura**, “What’s a Job Candidate Worth? Explaining Variation in Pay-Setting Practices,” *Working Paper*, 2020.
- Barach, Moshe A and John J Horton**, “How do employers use compensation history? Evidence from a field experiment,” *Journal of Labor Economics*, 2021, 39 (1), 193–218.
- Becker, Gary S**, *The economics of discrimination*, University of Chicago press, 1957.
- Blackaby, David, Alison L Booth, and Jeff Frank**, “Outside offers and the gender pay gap,” *The Economic Journal*, 2005, 115 (501), F81–F107.
- Blau, Francine D, Peter Brummund, and Albert Yung-Hsu Liu**, “Trends in occupational segregation by gender 1970–2009: Adjusting for the impact of changes in the occupational coding system,” *Demography*, 2013, 50 (2), 471–492.
- Cowgill, Bo, Amanda Agan, and Laura Gee**, “The Gender Disclosure Gap: Salary History Bans Unravel When Men Volunteer their Income,” *Working Paper*, 2022.
- Dee, Thomas S**, “A teacher like me: Does race, ethnicity, or gender matter?,” *American Economic Review*, 2005, 95 (2), 158–165.
- Faberman, R Jason, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa**, “Job search behavior among the employed and non-employed,” *Econometrica*, 2022, 90 (4), 1743–1779.
- Giuliano, Laura, David I Levine, and Jonathan Leonard**, “Manager race and the race of new hires,” *Journal of Labor Economics*, 2009, 27 (4), 589–631.
- Glassdoor**, “Progress on the Gender Pay Gap: 2019,” *White Paper*, 2019.
- Goldin, Claudia, Sari Pekkala Kerr, Claudia Olivetti, and Erling Barth**, “The expanding gender earnings gap: Evidence from the LEHD-2000 Census,” *American Economic Review*, 2017, 107 (5), 110–14.
- Krueger, Alan B**, “How computers have changed the wage structure: evidence from microdata, 1984–1989,” *The Quarterly Journal of Economics*, 1993, 108 (1), 33–60.
- List, John A**, “Non est disputandum de generalizability? A glimpse into the external validity trial,” Technical Report, NBER 2020.

- Mason, Malia F, Alice J Lee, Elizabeth A Wiley, and Daniel R Ames**, "Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations," *Journal of Experimental Social Psychology*, 2013, 49 (4), 759–763.
- Oehlert, Gary W**, "A note on the delta method," *The American Statistician*, 1992, 46 (1), 27–29.
- Rice, John A**, *Mathematical statistics and data analysis*, Cengage Learning, 2006.
- Roussille, Nina**, "The central role of the ask gap in gender pay inequality," *Quarterly Journal of Economics*, forthcoming.