

SALARY HISTORY AND EMPLOYER DEMAND:
EVIDENCE FROM A TWO-SIDED AUDIT

Amanda Y. Agan
Bo Cowgill
Laura K. Gee

WORKING PAPER 29460

NBER WORKING PAPER SERIES

SALARY HISTORY AND EMPLOYER DEMAND:
EVIDENCE FROM A TWO-SIDED AUDIT

Amanda Y. Agan
Bo Cowgill
Laura K. Gee

Working Paper 29460
<http://www.nber.org/papers/w29460>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2021, Revised July 2023

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 Roussille, 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, NBER Summer Institute (Labor), 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, Nuffield College (Oxford), RAND, Strategy Research Forum (SRF), Simon Fraser, SUNY Albany, University of California Santa Barbara, University of California Santa Cruz, University of Connecticut, University of Missouri, University of Pennsylvania Law School, University of Texas at Dallas and Urbana Champagne (UIUC). We also thank Hailey Brace, Karishma Chouhan, Nadine Fares, Zachary Finn, Matt Fondy and Norman Yuan for research assistance. And 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Amanda Y. Agan, Bo Cowgill, and Laura K. Gee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Salary History and Employer Demand: Evidence from a Two-Sided Audit
Amanda Y. Agan, Bo Cowgill, and Laura K. Gee
NBER Working Paper No. 29460
November 2021, Revised July 2023
JEL No. C90,J70,M50

ABSTRACT

We study how salary disclosures affect employer demand using a field experiment featuring hundreds of recruiters and over 2,000 job applications. We randomize the presence of salary questions and the candidates' disclosures. Employers make negative inferences about non-disclosing candidates, and view salary history as a stronger signal about competing options than worker quality. Disclosures by men (and other highly-paid candidates) yield higher salary offers, but are negative signals of value (net of salary), yielding fewer callbacks. Male wage premiums are regarded as a weaker signal of quality than other wage premiums (such as working at higher paying firms).

Amanda Y. Agan
Department of Economics
Rutgers University
75 Hamilton Street
New Brunswick, NJ 08901
and NBER
amanda.agan@rutgers.edu

Laura K. Gee
177 College Ave
Tufts University
Joyce Cummings Center
Medford, MA 02144
laura.gee@tufts.edu

Bo Cowgill
Columbia Business School
New York, New York
bo.cowgill.work@gmail.com

A data appendix is available at <http://www.nber.org/data-appendix/w29460>

1 Introduction

Labor markets are rife with asymmetric information. Knowing previous salaries can reduce uncertainty, influence employment decisions by firms, and shape income for workers. However, the meaning of disclosures can vary. In several important economic models,¹ workers are primarily differentiated by their outside options (which are unobservable to employers). In this setup, salary information could signal workers' outside options and bargaining positions. By contrast, an additional set of influential models features workers differentiated by "hidden ability."² Talent is often hard to measure, and salary information could reveal latent aptitude.

In practice, what information does a worker's salary reveal? What message is sent by not disclosing anything at all? The answers to 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 to study these questions. In our main empirical results, we execute a novel, two-sided field experiment in a real world setting. 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 opening were fictitious.

On the employer side, we randomly vary the presence of a salary history prompt on the job application form; and, on the candidate side, we randomly vary whether candidates disclose their current salaries or withhold them. Among candidates who disclose, we also vary the levels of salaries, with an average male-favoring gender wage gap. We then measure recruiters' callback recommendations and salary offers for all candidates. Recruiters are asked for a maximum salary offer (e.g., willingness-to-pay, or WTP) for all candidates, and about their beliefs about each worker's competing offers.

Our field experiment yields three main results. First, employers presume that "silent" candidates who do not disclose a salary have lower-than-average "hidden ability" as well as lower outside options. Unsurprisingly, these silent candidates are also given lower salary offers. Candidates' disclosures have a large impact on how recruiters assess their compet-

¹For example, models of monopsony (Manning, 2003; Card et al., 2018), efficiency wages (Lazear et al., 2016) and some models of on-the-job search (Burdett, 1978).

²For models with private worker ability see Gibbons and Murphy (1992) and Oyer and Schaefer (2011), and papers about market-based tournaments (Waldman, 2013; DeVaro and Kauhanen, 2016).

ing offers. Disclosures have a smaller impact on how recruiters assess candidates' quality (measured by WTP) — but both are affected in the same direction. Women are punished less for silence about their salary. Our supplemental workforce survey³ suggests a potential reason: women at all salary levels are less comfortable disclosing their salary, which makes their choice to disclose less revealing.

Our second set of results concerns the *amounts* of the salaries disclosed. On average, higher salaries raise employer beliefs about the worker's latent quality and competing offers, but beliefs about competing offers are affected more. Every \$1.00 increase in disclosed salary increases i) beliefs about the median competing offer by \$0.77, ii) employer WTP by \$0.65, and iii) the salary offer by \$0.68.

Salary variation arises from a variety of sources. Some candidates are employed at firms that are high (or low) paying, for all workers. Some candidates enjoy a male premium from gender wage gaps. Some candidates may have a positive wage residual (i.e., a higher salary even accounting for their firm, gender and other observables). Our experiment is designed to study how each type of salary variation affects candidates' outcomes.

We find that extra dollars given to men (through the gender wage gap) increase beliefs about both "hidden talent" for men and their outside offers. However, the gender gap differences in disclosure amounts have a bigger effect on assessments of outside options by gender. An extra dollar given to a male candidate (through the gender wage gap) raises WTP by only \$0.42, but increases beliefs about men's outside options by \$0.62. Male salary offers increase by \$0.48 for every extra dollar disclosed from the gender gap.

These effects are significantly larger than zero. Recruiters are discounting the male premium by about half, but far from 100%. However, the effects are significantly smaller than those from other sources of salary variation. Recruiters appear to *anticipate* overpaid men. By contrast, recruiters act as if extra dollars coming from other sources (e.g., working at a high-wage firm, or being well-paid within a firm's internal distribution) are far more informative. Each extra dollar from these sources increased WTP by \$0.64 to \$0.70 (more than \$0.20 more than the male gender gap bonus).

Together, our first two results contain common themes. Recruiters believe disclosures — both the choice to disclose and the amount — contain more information about a candidate's outside offers than a candidate's underlying quality. In addition, recruiters regard female silence (and male salary premiums) as less informative than other sources of variation. This may reflect recruiter awareness of true correlations between gender, compensation, value,

³We have written up the full details of our survey questions and results in a supplementary paper (Cowgill et al., 2022).

and willingness to disclose.

The final set of results are about how recruiters pick candidates and set salaries. As shown above, effects on salary offers are straightforward: disclosing workers — particularly those with high salaries — receive higher salary offers. Because our male candidates have higher salaries, this particularly benefits men.

By contrast, disclosing workers — especially those with high salaries — are less likely to be called back at all. Although they enjoy higher salary offers when selected, they are less likely to move forward in the hiring process. Our results about callbacks go in the opposite direction as those on salary amounts. Our interpretation is that callback decisions reflect expected costs. The idea that buyers lower the quantity demanded in response to higher prices appears in many economic models (e.g., the “law of demand”).

Our experiment connects this pattern to disclosures. Disclosures can suggest the price needed to outbid competitors is high. At some point, higher prices squeeze employers’ margins, which lowers the benefit of pursuing the candidate. Our experiment captures an estimate of employer margins, and results about callbacks move in parallel with these estimates.

Consistent with these reasons, silent candidates (and low-disclosers) are called back more often, but are offered lower average salaries when called. A recent non-experimental paper by [Kuhn et al. \(2022\)](#) finds similar results. This pattern has gendered effects: men are less likely to be chosen when they disclose. On the margin, employers interpret higher male salaries as a prohibitively high price tag. Women’s disclosures have a smaller effect on callbacks. We find similar results on the amounts of disclosures: lower amounts disclosed by women increase their odds of being recommended.

As a whole, our results suggest trade-offs. In our setting, disclosures — and disclosing higher amounts — increase the level of salary offers but lower callback rates. Additional compensation is obviously useful to candidates; however, additional outside offers can also give workers leverage and options. Implications for candidates’ disclosure strategies thus depend partly on how job-seekers and employers prioritize competing effects.

This paper provides four main contributions, which we detail in the next section. First, 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 extends the literature on “price as a signal of quality” in a labor market setting, and we develop the idea of “price as a signal of competition.” Second, we provide novel findings about gender differences in job search and how employers anticipate and react

to them. We propose a microfoundation for these differences in our setting, and we trace how these foundations affect candidate choices and employer reactions. In several empirical results, we find that recruiters incorporate expectations about gender differences into their decisions.

Finally, we develop a new experimental methodology to support the research questions around this topic. Our experimental design requires an extension of the audit methodology we call a “two-sided” audit. This gives us a novel, behind-the-scenes look at how information (and other interventions) propagate through hiring. We gather a rich collection of theoretically-motivated outcomes from the field about both wage setting and candidate selection. Our two-sided design creates multiple avenues for studying discrimination more broadly in future research.

The rest of the paper is organized as follows: Section 2 describes related literature and our contribution in more depth, and provides a brief background about the practice of asking job candidates for salary histories. Section 3 describes a conceptual framework of employer updating from salary information, and Section 4 describes our empirical setting. Section 5 lays out our experimental design and Section 6 proposes specifications. Our experimental results are in Section 7. Section 8 discusses the generalizability of our findings and implications for bans, and the final section concludes.

2 Background

2.1 Related Literature

Statistical Discrimination: Productivity and Outside Offers. A long literature studies how employers use observable characteristics to estimate hidden qualities. Classic statistical discrimination were motivated by “hidden talent” questions (Phelps, 1972; Arrow, 1973; Bohren et al., 2019, 2023). The recent literature about monopsony (Manning, 2003; Ashenfelter et al., 2010; Card, 2022) suggests another application of statistical discrimination: Using observable characteristics to predict each candidate’s outside options. If an employer can estimate the worker’s competing options, they could hire a strong worker at a low price.

The distinction has different implications: If wages mostly signal talent, then a high historical salary could be “good news” for prospective employers. If wages mostly signal outside options, then a high salary could be “bad news” to employers because it conveys a high price tag. We find evidence that salaries signal both talent and outside options, but that the effect

on outside options is generally bigger. In particular, our recruiters interpret gender-related salary differences as signals of male outside options, rather than of quality.

Voluntary Disclosure. The informational content of wages could also depend on how it was revealed. In many settings, workers face a disclosure choice. 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 dimension 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.

Gender differences in job search. 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, 2019](#); [Murciano-Goroff, 2017](#)), ask for lower salaries from employers ([Roussille, 2020](#)), and are less willing to disclose salaries ([Goldfarb and Tucker, 2012](#); [Cowgill et al., 2022](#)).⁴

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.⁵ 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

⁴Other 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, 2017](#)), 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](#)).

⁵Although 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 function directly. [Exley et al. \(2020\)](#) similarly studies negotiation costs by gender, including indirect costs of unsuccessful negotiations.

unraveling.

We then measure these responses empirically in a field experiment. We show novel evidence on employers' response to gendered negotiation behavior. Our results suggest that recruiters indeed interpret disclosure differently across genders, anticipating less disclosure from women, and punishing them *less* for silence. We also find that recruiters discount the higher salaries disclosed by men. Although recruiters' anticipation is insufficient to fully eliminate gender wage gaps, it does reduce them significantly below the original levels in our experiment.⁶

Our paper sheds new light on how employers interpret gendered differences in self-promotion and negotiation aggression. These are not only interpreted as a signal of candidate quality, but also as signals about the candidate's bargaining position and competitive alternatives. Less aggressive job search (such as non-disclosure) affects women twice: once through the employer's own assessment of the candidate's quality, and again through the employer's beliefs about how rival employers view the candidate.

Price as a signal of quality and/or competition. Our paper is also related to prior literature about using price as a signal of quality. This concept originally appeared in industrial organization studies of consumer products, but is less-developed in labor settings (Roussille, 2020). Seminal papers by Wolinsky (1983) and Milgrom and Roberts (1986), study price as a signal of quality, but do not portray price as a signal of competing offers, possibly because of the presumed thickness/competitiveness of demand for consumer products.⁷

By contrast, labor markets often feature thin and/or monopsonistic demand for workers (Manning, 2003; Ashenfelter et al., 2010). In this setting, a worker's price can signal not only their quality, but also the depth and quality of competing offers. These beliefs can affect wages through a separate, non-quality channel. We formalize this notion in our model and relate it to wage-setting in imperfectly competitive markets. The interaction of these signals impacts candidates' choices to reveal (or conceal) historical prices.

We then provide direct experimental evidence on how historical prices affect employer beliefs. In most of our results, we find a greater role for price as a signal of competition (versus as a signal of quality). Recruiters particularly interpret the gender-related salary

⁶Murciano-Goroff (2017) says that evidence about anticipation is "lacking" (p. 3). The most closely related papers regarding anticipation are Reuben et al. (2014), Exley and Kessler (2019) and Murciano-Goroff (2017). These papers report limited evidence of anticipation using laboratory (Reuben et al., 2014), online subject pools (Exley and Kessler, 2019), and observational (Murciano-Goroff, 2017) designs. Our data does not allow us to compare the degrees of anticipation across these papers.

⁷An unpublished manuscript by Allon et al. (2012) ("Price as a signal of availability") comes closest to developing this idea in the industrial organization setting.

differences as reflecting differences in candidates' outside offers, rather than differences in their "hidden abilities" on the job.

Audit Methodology. Methodologically, our work is related to recent innovations in correspondence audit methodology (Bartos et al., 2016; Kessler et al., 2019; Avivi et al., 2021). A review by Bertrand and Duflo (2017) says, "With a few exceptions, the literature has failed to push the [audit] correspondence methodology to design approaches to more formally test for various theories of why differential treatment is taking place." Our two-sided design allows researchers to collect detailed outcome data (beyond the binary callback choice) that reveal and suggest 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.

Information-Seeking Bans. Finally, our work relates to policies to "blind" decision makers to the personal history and identity of job applicants. Recent laws limit credit checks, drug tests, gender questions, criminal history questions and other personal details.⁸ In many settings, blinding requires cooperation from the supply side, which can override blinding through voluntary disclosure. Unraveling could unblind the decision (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 study salary history. A nascent literature directly studies these bans. The theoretical predictions about the effects of salary history bans are nuanced and ambiguous (Cullen and Pakzad-Hurson, 2021; Meli and Spindler, 2019). A series of empirical papers study salary history bans using panel methods and a variety of observational data sets (Bessen et al., 2020; Davis et al., 2020; Hansen and McNichols, 2020; Mask, 2020; 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., 2022).

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 competition). Our results address design considerations in policies for blinding decision-making. As the next section shows, there is significant variation in the design of

⁸Bartik and Nelson (2016); Friedberg et al. (2017); Corbae and Glover (2018) study credit checks, Card et al. (2021); Kuhn and Shen (2021) study gender information, and Doleac and Hansen (2020); Agan and Starr (2018) study criminal history questions.

salary history bans across jurisdictions.

2.2 Salary History Questions

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 evidence⁹ 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. 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.¹⁰ 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.¹¹ 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.¹²

The popularity of these bans masks enormous heterogeneity in their designs. Our findings in Section 8 address this heterogeneity. Some jurisdictions ban employers from asking **until an initial offer has been made** (but *can ask afterwards during negotiation of the initial offer*).¹³ By contrast, other jurisdictions ban asking **at any point in hiring**, both when making callback decisions and when setting wages.¹⁴ Our discussion in Section 8 connects our experimental results to the choice between these designs.

⁹See Cowgill et al. (2022).

¹⁰An incomplete list includes the states of New York, New Jersey, North Carolina, Vermont and Virginia, as well as municipalities including San Francisco, New Orleans, Kansas City, Atlanta, San Francisco, and others. See <https://www.hrdiver.com/news/salary-history-ban-states-list/516662/>.

¹¹See <https://www.hrdiver.com/news/salary-history-ban-states-list/516662/> for the most up-to-date list.

¹²Some jurisdictions' bans explicitly grant permission to confirm or use voluntarily supplied salary history information. One exception is California, where employers are expressly prohibited from relying on even voluntarily disclosed information.

¹³This design has been adopted in New Jersey, Alabama, Delaware, District of Columbia, New York (2017-2020, until a revision occurred), and Atlanta. See <https://www.hrdiver.com/news/salary-history-ban-states-list/516662/> and <https://www.ebglaw.com/news/new-york-state-releases-guidance-on-salary-history-ban/>.

¹⁴This includes California, and New York (from 2020 to the present, following a revision).

3 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. η is privately known to the worker, but the employer 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:

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

Information. The employer has a joint distribution of beliefs about the candidate’s value, outside offers 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 this relationships affects our theoretical predictions. Finally, our experiment measures this relationship empirically in a natural setting.

3.1 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 offers (given a disclosure action $r \in \{h, \emptyset\}$), the employer can calculate a TIOLI (take-it-or-leave-it) offer s^* :

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

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 v increase the employer's returns from giving a callback and for a generous TIOLI amount. However, higher beliefs about outside options η increase s^* , but decrease returns of sending a callback (unless beliefs about v simultaneously increase).¹⁵

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.

3.2 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 unravelling are the disclosure costs m . Because of these, silent workers contain a mixture of candidates with low salaries (strategically withheld), and candidates with high costs.

¹⁵By 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 suggests that their silence contains less information.

3.3 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 offer 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.¹⁶ 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 offers. However, it could be very noisy or informative. 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.¹⁷ Elite athletes are sometimes said to be “underpaid” because salary caps prevent stars from seeking their full value.¹⁸ 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 Case 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).

¹⁶For example, Manning (2003) reviews a series of “new classical monopsony.”

¹⁷In 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, 2017). This would potentially leave more buyer surplus on the higher end of the market.

¹⁸<https://www.npr.org/transcripts/628132840>

4 Empirical Setting

Our framework shows that the effect of disclosures is ambiguous and depends on specific characteristics of the setting. We now study how disclosure affects outcomes in a specific applied setting. We focus on the software engineering industry, in part because they help place our results in the context of wider labor market trends (Appendix A). In a later section (8), we discuss the generalizability of results from this setting.

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, 2021). Firms either hire individual recruiters on a temporary, contract basis, or they outsource recruiting entirely to a third party organization. In this section, we review the details of this practice. Although we relegate a discussion of the external validity to a later section, the section below will familiarize readers with this practice and demonstrate that our implementation choices reflect commonly-adopted practice.

4.1 Scope of Recruiter 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.¹⁹ Even before the COVID-19 pandemic, over 80% of outsourced recruiting was performed remotely.²⁰ Recruiters in this industry are often told to avoid searching on the Internet for information on job candidates, as this can violate employment law.²¹ Our study design closely mimics each of these attributes.

¹⁹See https://staging.kornferry.com/media/sidebar_downloads/Measuring-Up-A-new-research-report-about-RPO-metrics.pdf, <https://www.shrm.org/ResourcesAndTools/business-solutions/Documents/Talent-Acquisition-Report-All-Industries-All-FTEs.pdf>, and Analysts (2017) and our own survey responses from subjects in this experiment.

²⁰*Staffing Industry Analysts*, RPO Market Developments, December 2017.

²¹For example the Equal Employment Opportunity Commission (EEOC) tells firms to avoid online searching for candidates. See <https://www.lexology.com/library/detail.aspx?g=1147c039-ef9c-4f6a-9ebb-448de20b8123>.

4.2 Recruiter Incentive Pay

Bonuses — including both discretionary and formulaic kinds — are widespread in recruiting. Recruiters are primarily encouraged to fill open positions.²² 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.²³

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.²⁴ 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.”²⁵ On her recruiting podcast, HR consultant Laurie Ruettimann observed:²⁶

[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. 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.”²⁷ Textbooks for HR

²²According to industry reports, 60% of performance pay is measured by the number of positions filled. See <https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/rewarding-recruiters-for-performance.aspx>

²³In 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. For example, see HR Magazine <https://www.shrm.org/hr-today/news/hr-magazine/pages/1103hirschman.aspx>.

²⁴<https://www.cnn.com/audio/podcasts/diversifying/episodes/d19faee4-0af7-4047-b914-ae7001185387>.

²⁵<https://www.audacy.com/star94atlanta/latest/star-exclusive-viral-recruiter-tells-all>

²⁶<https://laurieruettimann.com/recruiter-mercedes-johnson-story/>

²⁷In some parts of the recruiting industry offered a bonus as a percentage of the candidate’s starting salary

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).

5 Experimental Design

To examine the 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.²⁸

5.1 Our Recruiting Workforce

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). Each subject was assigned to one of the experimental conditions (described below) using the randomization procedure in Appendix D.1.

(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.

²⁸To our knowledge, a similar design has been deployed in only one other paper (a working paper by Cowgill and Perkowski, 2021), and we have substantially extended the design here.

5.2 Our Recruiting Task

Recruiters were given three 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. Below, we detail the items in the structured candidate evaluation. Recent surveys of recruiters by Jobvite,²⁹ Monster.com,³⁰ and Black et al. (2020) indicate that assessing candidates using structured evaluation is typical and expected to grow.

We describe recruiter choices using economics terms in this manuscript, but our feedback form (and all communications with subjects) used everyday language. A full copy of the structured evaluation can be seen in Appendix N.

5.2.1 Primary Assessments

Callback. Like a traditional audit study, we observe whether each recruiter recommends a candidate for a callback. Recruiters were allowed to suggest more than one callback (as many as deemed a good fit).³¹ We conceptualize this choice as the callback modeled in Equation 3 of our framework.

Salary Offers. Recruiters made 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.³² The offer and callback decisions were on the first page of our evaluation.

²⁹https://www.jobvite.com/wp-content/uploads/2015/09/jobvite_recruiter_nation_2015.pdf

³⁰<https://www.monster.com/about/a/monster-2018-state-of-recruiting-survey>

³¹This is common in high-tech labor settings featuring high demand for qualified workers.

³²Recruiters 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. 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.”

Willingness to Pay (WTP). Recruiters are sometimes asked to match 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 a first offer. As a result, our recruiters reported the maximum offer from the candidate that the firm should accept.³³ 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](#)).

Recruiters were given a general sense of these topics via the task description, but were not given the specific questions in advance. Each question above 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.³⁴ 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.

5.2.2 Additional Assessments After the questions above, recruiters made the assessments below.

Outside Offer Distributions. For each candidate, recruiters state TIOLI salary offers that the candidate would be highly likely to accept, highly *unlikely* to accept, and indifferent about accepting. These were specifically defined as salary offers the candidate would accept with 95% probability, 50% probability and 5% probability. We interpret these as the recruiter’s beliefs about the distribution of the candidate’s best outside offer — in our conceptual framework, η — at the 5th, 50th and 95th percentiles of this distribution.

Number and Sources of Competing Offers. Recruiters 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.”³⁵ Recruiters also

³³Recruiters 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.

³⁴If we later found an error in an item (e.g., a recruiter typed in letters in an item that required numbers), we allowed the recruiter to fix the mistake.

³⁵Prior research suggests that outside offers 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. \(2017\)](#) find that only 29.1% of employed workers who are looking for work receive at least one offer per

state whether competing job offers would come from the candidate’s own search efforts, or from rival employers’ search efforts.

Our analysis uses a few simple combinations of the variables above (employer surplus, range of outside offers, probability of accepting the TIOLI offer, and probability of accepting \times employer surplus). These are explained in context, but we also document their assembly in Appendix E.

5.2.2.1 Objectives. 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). We discussed these objectives in Section 4. 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.

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 included text about a bonus in simple, non-technical language outlining the goals above. Additional details 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.³⁶

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.

5.3 Experimental Manipulation

Our field experiment contains experimental manipulation on both the employer side *and* the candidate side.

month.

³⁶The bonus includes a measure of candidate quality (assessed by the employer) to allow recruiters to pay more attract a higher quality candidate without penalty. Because there were no actual candidates nor firms, they were based on simulated outcomes based on data from comparable settings.

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

5.3.2 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 disclosures.

Biographical Details. As stated in our theory section, we had two hypotheses about biographical details. The first is about gender. 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.³⁷

The second randomized biographical detail is about low- and high-wage current firms. To randomize this, we used a list of the 13 biggest employers of software engineers from Monster.com and Indeed.³⁸ This included substantial variation in median salaries for software engineers who recently graduated.³⁹

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 which we divided into four packets of eight candidates. Each packet contained 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 “biography fixed effects.”⁴⁰

³⁷The names were: Andrew, Christopher, Joshua, Tyler, Emily, Jessica, Samantha, and Sarah.

³⁸<https://www.monster.com/career-advice/article/top-tech-employers-job-listings> and <https://www.techrepublic.com/article/the-10-companies-hiring-more-software-engineers-than-anyone-else-in-silicon-valley/>.

³⁹According to PayScale.com, the highest was Oracle (median salary of \$126K), and the lowest was General Dynamics (median salary of \$73K).

⁴⁰A biography consists of the specific combination of name, gender, previous and current employers, job

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.”⁴¹ 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).

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.⁴² Candidates were assigned a salary 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.

To understand how disclosure affect 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 92% of the average male disclosed salary — however the actual wage gaps recruiters observe within their eight applications varied.⁴³ Figure 1 shows the resulting distribution of disclosed salaries amongst candidates who dis-

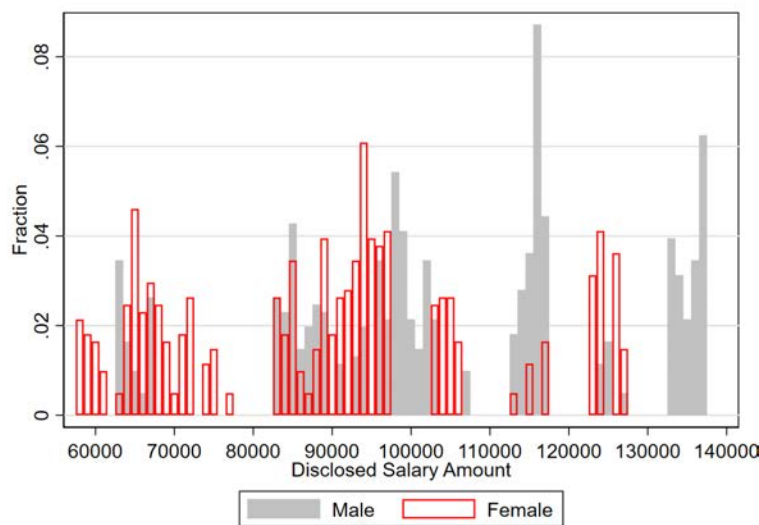
titles and descriptions, dates of employment, undergraduate education, and skills.

⁴¹This 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, however slightly.

⁴²Data from websites like Payscale.com and Glassdoor.com are self-reported by workers who visit these 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).

⁴³Amongst recruiters who saw any disclosed salaries, the gender gap they saw varied from females disclosing salaries that were on average 61% lower to 14% higher (at the 5th and 95th 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 gap in wages seen by all our recruiters is 92% 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 – do not contain gender-specific wage values.

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).

closed 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. We also add significant random variation in salaries, conditional on gender and prior employer, and uncorrelated with any other characteristic. Appendix F contains additional details about how we assigned salaries to candidates.

5.4 Subtreatments

We clustered randomization in three ways. First, we held the 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 clustered 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 clustering allowed us to measure how candidates’ disclosure choices affected each other through spillovers.

Finally, we clustered candidates’ disclosure amounts. As previously mentioned, candi-

dates' disclosure amounts were randomized (for those who disclosed). We clustered 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 clustered 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).

6 Specifications

Our conceptual framework shows how the effects of disclosure are ambiguous, both in direction and magnitude. To measure directions and effects in our empirical setting, we mainly use the three 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).

These tests measure the effects of disclosure behavior on final outcomes such as salary offers and callbacks, and additionally measure mechanisms (corresponding to those in our conceptual framework). For example, our theory highlights the way disclosures can impact both WTP and the employer's beliefs about the candidate's competing offers. Which of these is affected more impacts the incentives to disclose (and how disclosures are interpreted). As such, we test for differences of treatment effects across two outcome variables (WTP and outside offers),⁴⁴ and share the differences (and p -values) in the main text.

6.1 Specification 1: Disclosure vs. Silence

$$y_{ij} = \beta_1 \text{SalaryDisclosed}_{ij} + \beta_2 \text{SalaryHistoryAsked}_j + v_i + \beta_3 [\text{RecruiterControls}_j] + \epsilon_j \quad (4)$$

where i indexes candidates and j indexes recruiters. Whether the recruiter saw applications with a salary history question is controlled for by $\text{SalaryHistoryAsked}_j$. $\text{RecruiterControls}_{ij}$ includes the gender, race, experience level, and hourly rate of the assigned recruiter (all

⁴⁴These tests use seemingly unrelated regression.

balanced by design, section D.1). v_i signifies candidate biography fixed effects (see section 5.3.2).

We extend this specification in two main ways: First, we introduce candidate gender interactions with the $SalaryDisclosed_{ij}$ terms.⁴⁵ Second, we replace $SalaryDisclosed_{ij}$ with a vector $\{DisclosedLowSalary_{ij}, DisclosedHighSalary_{ij}\}$ where $DisclosedLowSalary_{ij}$ is a dummy for whether the applicant disclosed a salary at the 25th percentile of his or her current firm and $DisclosedHighSalary_{ij}$ was the same for the 75th percentile.

Interpretation. The $SalaryDisclosed_{ij}$ coefficient (β_1) measures whether disclosers are (on average) rewarded more or less than silent types. A positive coefficient indicates that recruiters believe positive self-selection into disclosing (i.e., above-average types choose to disclose). The 25th versus 75th percentile dummy variables help measure where silent candidates are presumed to fall in the overall distribution.

6.2 Specification 2: Effects of Disclosure Amounts

We now add $AmountDisclosed_{ij}$ terms to Equation 4.

$$\begin{aligned}
 y_{ij} = & \beta_1 SalaryDisclosed_{ij} + \beta_2 SalaryDisclosed_{ij} \times AmountDisclosed_{ij} \\
 & + \beta_3 SalaryHistoryAsked_j + v_i + \beta_4 [SpilloverControls_{ij}] \\
 & + \beta_5 [RecruiterControls_j] + \epsilon_j
 \end{aligned}
 \tag{5}$$

We set $AmountDisclosed = 0$ for candidates that did not disclose (their overall impacts are captured by “ $SalaryDisclosed = 0$ ”). Spillover controls are detailed in Appendix G.⁴⁶

Interpretation. Larger coefficients on $AmountDisclosed_{ij}$ (β_2) (steeper slopes) indicate greater impacts on a recruiter’s evaluation. Greater magnitudes indicate recruiters relying on the disclosure amount and updating beliefs about the candidate based on the disclosure amount. The opposite extreme (a low, flat $AmountDisclosed_{ij}$ coefficient) indicates recruiters do not change or incorporate new information from the disclosure amount. We will use this specification to measure the monotonicity assumption in our theoretical framework (Assumption

⁴⁵The main effect of candidates’ gender is absorbed by the biography fixed effects (v_i).

⁴⁶Our main interest in this paper is the average direct effects of one’s disclosures on one’s own outcomes, but as described in Section 5.4, 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.

2). This requires that the employer’s surplus be increasing (or decreasing) in salary history amounts.

6.3 Specification 3: Heterogeneous Effects of Disclosing an Additional \$1

Finally, we use a third specification about components of $AmountDisclosed_{ij}$. When we constructed candidate salaries (5.3.2), three factors went into candidate salaries: 1) Some work at higher or lower -wage firms, 2) some are relatively well- or poorly- paid within their firm’s distribution, and 3) some are male or female. Recruiters could plausibly treat each source of pay variation equally. 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. In order to measure these differences in our data, we use the following identity:

$$AmountDisclosed_{ij} = SampleAverage + FirmOffset_i + GenderOffset_i + WithinFirmOffset_{ij} \quad (6)$$

Each candidate’s disclosure amount is the sum of an overall sample average across all candidates ($SampleAverage$, a constant), plus a firm-specific offset for the candidate’s employer ($FirmOffset_i$, some firms pay higher or lower on average), a gender offset ($GenderOffset_i$, penalizing women and favoring men), and a within-firm offset ($WithinFirmOffset_{ij}$, representing pay variation for the same job). The sum of these is equal to the total amount disclosed. These relationships flow directly from our procedure for creating salary disclosure amounts (Section 5.3.2). We use the definition in Equation 6 to replace $AmountDisclosed_{ij}$ in Equation 5 with its subcomponents. This leads to our third specification:

$$\begin{aligned} y_{ij} = & \beta_1 SalaryDisclosed_{ij} + \beta_2 SalaryDisclosed_{ij} \times FirmOffset_i \\ & + \beta_3 SalaryDisclosed_{ij} \times GenderOffset_i \\ & + \beta_4 SalaryDisclosed_{ij} \times WithinFirmOffset_{ij} \\ & + \beta_5 SalaryHistoryAsked_j + v_i + \beta_6 [SpilloverControls_{ij}] + \beta_7 [RecruiterControls_j] + \epsilon_j \end{aligned} \quad (7)$$

This allows us to obtain separate coefficients for each source of salary variation, and compare them. We call these the heterogeneous effects of an additional \$1. Because prior employer, gender and within-firm salary are randomly assigned, the “salary offset” coefficients can each be interpreted causally. These effects are separately identified from the direct effects of being male and employed at a high-wage firm, because our experiment contains identical candidates (same genders and employers) who were silent about their salary history.

Interpretation. As with the slopes in the previous regression (5), the magnitude of these coefficients represents how informative they are. If recruiters anticipate any differences by using biographical features, this anticipation would *reduce* the informational content of the disclosure, and push these coefficients toward zero. If gender and employer differences are *not* fully anticipated, we would find nonzero results.

7 Results

In the sections below, we study how recruiters reacted to salary disclosure versus silence (7.1) and to higher versus lower salary amounts (7.2 and 7.3). In our last section of results (7.4), we study how disclosures affected choices about callbacks. Table 2 shows descriptive statistics from our experiment, which we discuss more below.

Table 1: Candidate Summary Statistics

	All			Male Candidate			Female Candidate		
	(1) All	(2) Salary Disclosed	(3) Not Disclosed	(4) All	(5) Salary Disclosed	(6) Not Disclosed	(7) All	(8) Salary Disclosed	(9) Not Disclosed
WTP	107,101	109,922	102,978	110,814	114,873	104,882	103,388	104,970	101,075
Outside Offer 5th %ile	88,836	93,772	81,622	92,664	99,727	82,342	85,007	87,816	80,902
Outside Offer 50th %ile	100,044	103,635	94,796	104,235	109,783	96,127	95,852	97,486	93,465
Outside Offer 95th %ile	112,161	113,546	110,136	117,517	119,697	114,332	106,804	107,395	105,940
Outside Offer Range	23,325	19,774	28,514	24,853	19,970	31,990	21,797	19,579	25,038
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
Surplus	6,144	5,929	6,458	6,226	5,766	6,899	6,061	6,091	6,016
p(accept)	0.546	0.536	0.560	0.528	0.506	0.561	0.563	0.565	0.560
p(accept) \times Surplus	3,268	3,200	3,367	3,133	2,881	3,500	3,403	3,519	3,234
≥ 2 Other Offers	0.530	0.526	0.536	0.560	0.549	0.575	0.501	0.503	0.498
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 5.2 and Appendix E.

7.1 Disclosure vs. Silence

Our results show that on average, our recruiters believe silent candidates have lower quality (measured by WTP) and lower outside options. We also find they are given lower salary offers (versus the average disclosing candidate). We can see this pattern in basic descriptive statistics in Table 1: recruiters' WTP, beliefs about outside offers, and TIOLI offers are lower for candidates who do not disclose (column 3) compared to those who do (column 2).

In Table 2 (odd columns), we measure these differences using our disclosure versus silence regression (Equation 4). When a candidate discloses, recruiters raise willingness-to-pay by about \$6,800 on average (or 6.5% over the mean of non-disclosers). Likewise, disclosure increases beliefs about the level of outside offers by \$8,400 on average (+8.8% over the mean for non-disclosers). In Column 7, we see that disclosure also reduces uncertainty about outside offers. Recruiters' beliefs about the distribution of acceptable offers are compressed by \$9,200 on average (or 32% over mean spread of \$28,500) between the 5th and 95th percentile.

Table 2: Average Effect of Disclosing Salary

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Offer 50th %ile	Outside Offer 50th %ile	Offer	Offer	Outside Offer Range	Outside Offer Range
Salary Disclosed	0.68*** (0.14)	0.97*** (0.16)	0.84*** (0.14)	1.30*** (0.15)	0.73*** (0.14)	1.08*** (0.16)	-0.92*** (0.24)	-1.25** (0.40)
Female x Disclosed		-0.57*** (0.13)		-0.92*** (0.11)		-0.69*** (0.12)		0.67+ (0.37)
Female Disclosure Effect:								
<i>Total</i>		0.40		0.38		0.39		-0.59
<i>p-value</i>		0.01		0.01		0.01		0.00
Mean Non-Disclosers:								
<i>All</i>	10.30	10.30	9.48	9.48	9.65	9.65	2.85	2.85
<i>Male</i>	10.49	10.49	9.61	9.61	9.80	9.80	3.20	3.20
<i>Female</i>	10.11	10.11	9.35	9.35	9.51	9.51	2.50	2.50
R ²	0.18	0.18	0.25	0.26	0.20	0.21	0.01	0.02
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: This table shows estimates from versions of Equation 4, including recruiter controls and candidate fixed effects. Dependent variables are listed in Section 5.2 and Appendix E. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Robust standard errors are clustered at the recruiter level. Table M4 uses the same specification to examine additional outcomes. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

These results suggest anticipation by recruiters, who appear to intuit positive selection into disclosing. Disclosure choices have a bigger effect on recruiters' beliefs about competing offers, and a smaller effect on willingness-to-pay. Testing the effect on WTP (Column 1) equals that on outside offers (Column 3) yields a p -value of 0.038. However, both are affected in the same direction. The salary offers given to disclosing candidates are about \$7,300 (7.5%) higher than silent ones. Our results suggest that outcomes for silent workers are worse than

the average candidate within the same {current employer \times job title \times gender} cell. Using the multiple hypothesis testing correction from [List et al. \(2019\)](#), we find that these outcomes remain statistically significant after the MHT correction.

The results above are average effects, but our design allows us to decompose the effect into low versus high earners (compared to firm averages). In [Table 3](#), we find that recruiters infer that silent candidates' hidden salaries are at (or just slightly below) the 25th percentile given the candidates' observables. Workers below this percentile are better off silent.

Table 3: Average Effect of Disclosing a High versus Low Salary

	(1)	(2)	(3)	(4)
	WTP	Outside Offer 50th %ile	Offer	Outside Offer Range
Disclosed 25th %ile Salary	-0.08 (0.14)	-0.04 (0.13)	-0.04 (0.12)	-0.52 (0.34)
Disclosed 75th %ile Salary	1.11*** (0.15)	1.24*** (0.14)	1.18*** (0.13)	-0.37 (0.29)
Mean Non-Disclosers	10.30	9.48	9.65	2.85
R ²	0.34	0.46	0.38	0.01
Observations	2048	2048	2048	2048

Notes: This table contains the results of [Eq. 4](#) (including recruiter and spillover controls, candidate FEs and sub-treatment FEs). Dependent variables are listed in the column header and explained in [Section 5.2](#) and [Appendix E](#). Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Disclosed Xth %ile Salary means a candidate disclosed a salary at the Xth percentile within their specific firm. The omitted category is candidates who did not disclose a salary. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

“Below the 25th percentile” is a significant discount. However, the theoretical literature on disclosure rationalizes even worse. Why aren't recruiters more punitive? 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 survey results find supporting evidence for this: workers are less willing to disclose extremely high salaries.⁴⁷

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.

⁴⁷See [Cowgill et al. \(2022\)](#).

Gender and Disclosure Costs. In addition, some workers may be inherently uncomfortable disclosing. Revealing salary history may feel repugnant. Multiple studies suggests that women in particular do not like disclosing (Goldfarb and Tucker, 2012; Cowgill et al., 2022).⁴⁸ Our survey results suggest that women are less willing to disclose their salaries, even after controlling for salary, education and other characteristics. If recruiters anticipate this, they may find female silence less informative.

This is what we find in our recruiter experiment. In the even columns of Table 2, we interact disclosing with gender. We find that women job candidates are punished *less* for non-disclosure than men. Recruiters penalize a silent man’s WTP by \$5,700 more, his outside option \$9,200 more, and his offer by \$6,900 more than for a silent woman. The flip side of this behavior, however, is that the benefit of disclosing is also smaller for women.

A potential alternative explanation for our results is that recruiters simply misjudged the average market wages for this job. We explore this in Appendix I. Lack of knowledge does not seem to be driving our results.

7.2 Amounts Disclosed

We now measure how recruiters reacted to higher or lower disclosed salaries in the odd columns of Table 4 (regression 5). On average, higher salaries increase recruiter WTP for candidates. 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 \$0.65 (column 1).

Column 3 shows that recruiters’ expectations about outside offers similarly increase with higher amounts, and is actually more responsive to disclosure amounts than WTP. For every \$1 of current salary disclosed, this increases by \$0.77 (compared to \$0.65 for WTP, $p = 0.002$).

In the even columns of Table 4, we also examine whether amounts disclosed have a different effect for male and female candidates. In aggregate, do recruiters respond to an extra dollar disclosed from men and women similarly? For most of the outcomes in Table 4, we do not find evidence that extra total dollars disclosed by female candidates are treated differently from those disclosed by male candidates. The coefficients on “Female \times Disclosed \times Amt Disclosed” are all statistically insignificant and small. However, these results are highly aggregated, and we decompose salary variation into multiple sources using Equation 5 next.

⁴⁸In our own survey, we find that women are about 12 percentage points less likely to disclose than men when not prompted, and that women disclose about 1 percentage point less often than men when prompted (Appendix Table K1). The difference when prompted is not statistically significant.

Table 4: Average Effect of Disclosing by Salary Amount

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Offer 50th %ile	Outside Offer 50th %ile	Offer	Offer	Outside Offer Range	Outside Offer Range
Disclosed x Amt Disclosed	0.65*** (0.06)	0.67*** (0.06)	0.77*** (0.05)	0.78*** (0.05)	0.68*** (0.05)	0.69*** (0.06)	-0.03 (0.12)	0.03 (0.18)
Salary Disclosed	-5.76*** (0.56)	-6.17*** (0.61)	-6.82*** (0.49)	-7.11*** (0.54)	-6.04*** (0.53)	-6.33*** (0.56)	-0.19 (0.92)	-1.12 (1.68)
Female x Disclosed x Amt Disclosed		0.01 (0.05)		-0.00 (0.04)		0.02 (0.04)		-0.03 (0.14)
Female x Disclosed		0.34 (0.47)		0.27 (0.39)		0.19 (0.41)		1.02 (1.49)
Female Amount Disclosed Slope:								
<i>Total</i>		0.68		0.78		0.71		-0.00
<i>p-value</i>		0.00		0.00		0.00		0.99
Mean Non-Disclosers:								
<i>All</i>	10.30	10.30	9.48	9.48	9.65	9.65	2.85	2.85
<i>Male</i>	10.49	10.49	9.61	9.61	9.80	9.80	3.20	3.20
<i>Female</i>	10.11	10.11	9.35	9.35	9.51	9.51	2.50	2.50
R ²	0.38	0.38	0.54	0.54	0.44	0.44	0.01	0.01
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: This table shows estimates from versions of Equation 5 (including interactions with gender, as well as recruiter controls, spillover controls, candidate FEs and sub-treatment FEs). Dependent variables are listed in the column header and explained in Section 5.2 and Appendix E. Salary amounts and outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Robust standard errors are clustered at the recruiter level. Table M5 uses the same specification to examine additional outcomes. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

7.3 Heterogeneous Effects of an Extra \$1

Variation in current salaries can arise from multiple sources. Table 5 decomposes $AmountDisclosed_{ij}$ into multiple sources of salary variation (specification 7). We broadly find that not all dollars are created equally. Within-firm variation —labeled “+\$10K within Firm” in Table 5— is the most informative to recruiters in the sense that it has the steepest slope: each additional \$1 disclosed from this within firm variation increases WTP by \$0.70. In theory, this variation is the only true source of surprise to the recruiter, as it is uncorrelated with anything else on the application form.

By contrast, the extra dollars from the gender wage premium for men, or from working at a high-wage firm, did not increase WTP by as much (these are represented as “+\$10K from Firm” and “+\$10K from Male,” respectively, in Table 5).⁴⁹ This is consistent with recruiters anticipating that a worker at Apple (a high-wage firm in our sample) is paid around a certain average, and that men are typically paid more. Although these slopes are generally lower than the unexpected “+\$10K from within Firm” surprise coefficient, our results suggest that recruiters do not fully anticipate wage differences at the prior firm.

⁴⁹We can statistically reject the difference with the gender wage premium ($p < 0.01$), but not with the high-

The gender slope is the flattest and least informative to the recruiters: for each additional \$1 candidates disclose because of the gender wage gap which favors men, recruiters give salary offers that are \$0.48 higher. Our results show that compared to other sources of variation, recruiters discount extra dollars given to men. They may interpret that men are paid extra for spurious reasons and thus find their higher salaries less informative (both about quality, and about the candidate’s outside offers). Alternatively, recruiters may feel that men’s higher reported salaries are uninformative because they are exaggerated.

Gender, WTP, and Outside Options. An extra \$1 from the gender wage gap is particularly uninformative about worker quality. For every \$1 given for random within-firm reasons (i.e., one of the non-gender reasons), recruiters increased WTP by \$0.70 (a high amount). However, for every \$1 allocated to men through the gender wage gap, recruiters updated WTP only by \$0.42 (much lower, $p < 0.01$).

wage firm premium ($p = 0.27$).

Table 5: Effect of An Extra Dollar Decomposed

	(1)	(2)	(3)	(4)	(5)
	WTP	Outside Offer 50th %ile	Offer	O.O. Range	p WTP = Outside Offer
+\$10k from Firm	0.64*** (0.07)	0.77*** (0.06)	0.68*** (0.06)	-0.01 (0.20)	0.00
+\$10k from Male	0.42*** (0.09)	0.62*** (0.07)	0.48*** (0.09)	-0.39 (0.26)	0.00
+\$10k within Firm	0.70*** (0.06)	0.79*** (0.05)	0.73*** (0.05)	0.04 (0.06)	0.06
p F-M	0.00	0.01	0.00	0.12	
p F-W	0.27	0.69	0.28	0.77	
p M-W	0.00	0.01	0.00	0.09	
R ²	0.38	0.54	0.44	0.01	
Observations	2048	2048	2048	2048	

Notes: This table shows estimates from Equation 7 (including recruiter and spillover controls, candidate FEs and sub-treatment FEs). Dependent variables are described in Section 5.2 and Appendix E. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. Column 5 reports the p -value on a test of whether the coefficient on WTP in Column (1) = the coefficient on Outside Offer 50th %ile in Column (2). p -values for comparisons of coefficients within the same model are provided in the 2nd panel. p F-M is the p -value from the test of “+\$1 from Firm” = equals “+\$1 from Male”. Robust standard errors are clustered at the recruiter level. Table M6 uses the same specification to examine additional outcomes. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

By contrast, recruiters were more willing to update beliefs about outside offers in response to gender wage differences. Each extra dollar given to men in our experiment for gender reasons increased beliefs about the median offer by \$0.62 (higher than the \$0.42 for WTP, $p < 0.01$). Together, these results suggest recruiters are reluctant to believe that men’s higher salaries are a signal of higher quality for men. However, recruiters are more willing to believe that higher salaries are a signal of higher competing offers for men (versus women). Although the inferences made from gender wage differences are smaller, they are still far from zero, and still impact recruiters’ assessments.

The pattern above holds across our study more generally: Our recruiters update beliefs on outside options using salary history more than beliefs about worker quality (WTP). When we examine the other two components of salary variation (between- and within- firm), we find the same pattern ($p < 0.01$). Higher outside offers could come partly from a candidate’s search effort, or from rival companies’ search efforts targeting the candidate. Our results suggest that recruiters believe that men enjoy higher outside options because of more candidate-driven search effort (Appendix Table M6).

7.4 Decision-Making Outcomes: Whom to Hire and How to Pay

Our final set of results is about how recruiters synthesize inferences into decisions. These decisions are modeled theoretically in Equations 2 (salary amounts) and 3 (callbacks). As we covered above, the effects on salary offers are straightforward: disclosing workers—particularly those with high salaries—receive higher salary offers.

Our results about callbacks are more nuanced. On average, our results about callbacks generally go in the opposite direction as those on salary offers. Our evidence suggests that disclosing workers—especially men and workers with high salaries—are less likely to be recommended at all. 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 framework (Equation 3), this happens because higher prices diminish employer surplus from hiring a worker.


Our estimates about callbacks are sometimes underpowered because of the binary nature of the outcome variable. In addition, the relationship between callbacks and amounts disclosed is noisy below \$70K. This may have arisen because our instructions informed recruiters that most workers in this role make between \$70K and \$120K (although we welcomed them to suggest something else).

To shed light on these results, we supplement our analysis in this section with outcomes

that reveal patterns and mechanisms. Our results in previous sections show that salary disclosures increase beliefs about competing offers, and thus raise the price necessary to outbid a rival employer. To better understand our callback results, we examine our experimental outcomes about employer surplus and probability of acceptance.

7.4.1 Callbacks and Amounts Our strongest evidence about callbacks comes from our analysis of amounts. Across several specifications, disclosing a higher salary reduces callbacks. Even in our full sample, which includes the noisy results on the low salaries, the relationship is downward sloping (albeit statistically insignificant, Table 6). When we focus on the sample above \$70K (the upper 85% of the data) in Table 7, we find the stronger negative relationships.⁵⁰ Figure 2 contains a visualization.

Figure 2: Callbacks and Amounts Disclosed



Notes: In this figure the circles present the proportion of job candidates who were recommended for a callback by the amount they disclosed. There is a line of best fit for observations below \$70K, and one for above \$70K (the upper 85% of the data).

Tables 6 and 7 present our results as regressions. These tables include outcomes that are theoretically-related to callbacks (employer surplus, predicted probability that candidate accepts offer, and probability of acceptance \times surplus) for both the full and truncated samples.

⁵⁰To better understand our data, we ran regressions predicting callback using Equation 5, both as is and with an additional term (square of salary). We ran these specifications for subsets for the data from those above 60K, 70K, 80K and so on. We found the coefficient on the amount disclosed is negative and statistically significant for amounts \$70K and above. The squared term is insignificant after we subset to 70K and above. Additional results are available upon request.

Every \$10,000 extra disclosed by a candidate reduces the probability of a callback by 4 percentage points (6.25% over the mean for non-disclosing candidates, Table 7). Table 6 shows that higher salaries not only reduce callbacks, but also the key drivers of the callback decision. Every \$1 extra disclosed by a candidate reduces employer surplus by \$0.03. Higher disclosure amounts also reduce beliefs about the candidate accepting the offer. Table 7 also presents gender interactions with amounts. We find that the amount/callback relationships in column 2 are generally stronger for men.

Table 6: Average Effect of Disclosing by Salary Amount for Callback Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Callback	Callback	Surplus	Surplus	p(accept)	p(accept)	p(accept) x Surplus	p(accept) x Surplus
Salary Disclosed	0.13 (0.14)	0.18 (0.18)	0.28 (0.17)	0.17 (0.21)	0.15 (0.12)	0.12 (0.14)	0.33* (0.13)	0.26+ (0.15)
Female x Disclosed		-0.17 (0.17)		0.15 (0.19)		0.03 (0.10)		0.10 (0.12)
Disclosed x Amt Disclosed	-0.01 (0.01)	-0.02 (0.02)	-0.03* (0.02)	-0.03 (0.02)	-0.02 (0.01)	-0.02 (0.01)	-0.03** (0.01)	-0.03* (0.01)
Female x Disclosed x Amt Disclosed		0.02 (0.02)		-0.01 (0.02)		-0.00 (0.01)		-0.01 (0.01)
Female Amount Disclosed Slope:								
<i>Total</i>		0.01		-0.04		-0.02		-0.04
<i>p-value</i>		0.71		0.07		0.23		0.01
Mean Non-Disclosers:								
<i>All</i>	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
<i>Male</i>	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
<i>Female</i>	0.62	0.62	0.60	0.60	0.56	0.56	0.32	0.32
R ²	0.04	0.04	0.04	0.04	0.07	0.07	0.06	0.06
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: This table shows estimates from versions of Equation 5 (including recruiter controls, spillover controls, candidate FEs and sub-treatment FEs). Dependent variables are listed in the column header. Salary amounts and outcomes measured in dollars (e.g. Surplus) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

7.4.2 Heterogeneous Effects of +\$1 As we showed above, we can disaggregate the effects of salary differences. Table 8 examines callback-related outcomes using the full sample (Equation 5). This allows us to measure the effects of disclosing \$1 more on callback outcomes, depending on whether the dollar came from the gender wage gap, between-firm variation or within-firm variation. For nearly all of our callback-related outcomes, we find negative relationships with extra dollars. Randomly assigned higher amounts from any source we study are correlated with lower callback outcomes, although some of these effects are statistically insignificant.

Of these, our results on the gender wage gap (\$10K from being male) are especially pre-

Table 7: Average Effect of Disclosing by Salary Amount for Callback Outcomes Salary > \$70,000

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Callback	Callback	Surplus	Surplus	p(accept)	p(accept)	p(accept) x Surplus	p(accept) x Surplus
Salary Disclosed	0.43*	0.56*	0.22	0.29	0.30+	0.28	0.37*	0.42*
	(0.18)	(0.22)	(0.21)	(0.26)	(0.16)	(0.17)	(0.15)	(0.16)
Female x Disclosed		-0.34		-0.19		0.03		-0.09
		(0.21)		(0.24)		(0.13)		(0.14)
Disclosed x Amt Disclosed	-0.04*	-0.05*	-0.03	-0.04	-0.03+	-0.03+	-0.04**	-0.04**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)
Female x Disclosed x Amt Disclosed		0.04+		0.02		-0.00		0.01
		(0.02)		(0.02)		(0.01)		(0.01)
Female Amount Disclosed Slope:								
<i>Total</i>		-0.01		-0.02		-0.03		-0.03
<i>p-value</i>		0.46		0.44		0.08		0.04
Mean Non-Disclosers:								
<i>All</i>	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
<i>Male</i>	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
<i>Female</i>	0.62	0.62	0.60	0.60	0.56	0.56	0.32	0.32
R ²	0.04	0.04	0.04	0.04	0.07	0.07	0.04	0.04
Observations	1849	1849	1849	1849	1849	1849	1849	1849

Notes: This table shows identical specifications from Table 6 (Equation 5, including recruiter controls, spillover controls, candidate FEs and sub-treatment FEs), restricting the sample to those that did not disclose a salary and those that disclosed a salary above \$70,000. Dependent variables are listed in the column header. Salary amounts and outcomes measured in dollars (e.g. Surplus) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

cisely estimated. We find that \$10,000 extra salary given to men because of the gender wage gap causes a 5 percentage point *decrease* in the probability of a callback, and reduces expected employer surplus by \$7,000. These results suggest that the higher wages afforded to men (through the gender wage gap) harm their callback chances when disclosed.

Men’s higher salaries appear to lower their salary-net value as candidates, in some cases by more than other sources of variation. This is the natural extension of the earlier findings. Throughout several tests, we find that salary history disclosure increases beliefs about outside offers more than candidate quality (Tables 2, 3 and 5). As outside offers outpace candidate quality, eventually employers’ margins could be squeezed. As salary expectations rise too high, selecting an expensive candidate may become a less lucrative deal.

Given these findings, some candidates may be better off silent. We examine this question in Appendix J. Our results on this question are sometimes imprecise, but lean closer to “silence increases callbacks” than otherwise, particularly for men and highly-paid candidates.

Table 8: Effect of An Extra Dollar Decomposed for Callback Outcomes

	(1)	(2)	(3)	(4)
	Callback	Surplus	p(accept)	p(accept) x Surplus
+\$10k from Firm	-0.02 (0.02)	-0.04+ (0.02)	-0.02 (0.01)	-0.04** (0.01)
+\$10k from Male	-0.05+ (0.03)	-0.07* (0.03)	-0.03 (0.02)	-0.05* (0.02)
+\$10k within Firm	0.00 (0.02)	-0.03 (0.02)	-0.01 (0.01)	-0.03+ (0.02)
p F-M	0.27	0.25	0.76	0.76
p F-W	0.28	0.68	0.46	0.36
p M-W	0.07	0.19	0.37	0.36
R ²	0.04	0.04	0.07	0.06
Observations	2048	2048	2048	2048

Notes: This table shows estimates from Equation 7 including recruiter and spillover controls, candidate FEs and sub-treatment FEs. Dependent variables are described in Section 5.2 and Appendix E. Outcomes measured in dollars (e.g. Surplus) are in \$10K increments. p -values for comparisons of coefficients within the same model are provided in the 2nd panel, where for example p F-M is the p -value testing that the coefficient from “+\$1 from Firm” = the coefficient on “+\$1 from Male”. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

8 Discussion

8.1 Generalizability of Results

Our conceptual framework provides intuition for drivers behind our results. However, the framework also shows how results could differ in other settings. In this section, we discuss the generalizability of our findings. 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. Appendix L discusses our findings using the SANS conditions (selection attrition, naturalness, and scalability) suggested by List (2020). We highlight two key aspects of this analysis:

- 1) **Attrition.** 14% of subjects did not complete the task. Recruiters assigned a packet that did *not* ask for salary history were more likely to drop out. In addition, recruiters assigned a packet where no single candidate provided their salary history (i.e., all were “silent”)

were more likely to drop out.

Our treatments were not designed to study recruiter retention. To be clear, this idea does not appear in any form within the conceptual framework guiding our thinking (Section 3). We discovered our attrition results only by checking standard attrition diagnostics. However, the behavioral forces producing these forces are easy to imagine. Without the guidance of salary histories, the recruiter's job could require more effort.

- 2) **Interviews.** According to our survey, the most common way to ask salary history is on the job application (in writing). However, we acknowledge that 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, 2020). 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.

Our empirical results most resemble Case 3 of our conceptual framework. They suggest that high h workers are better quality, but are also pursued aggressively by competing employers (which diminish 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.

We do not aim to generalize our results to the entire economy. To the contrary, our conceptual framework provides guidance about how our results depend on specific empirical properties and mediating factors of a given labor market (summarized in 3.3).

8.2 Implications for Bans

Salary disclosures are interesting partly because bans have been enacted. As we have discussed, a ban on questions does not prevent voluntary disclosure. A key question is therefore whether recruiters interpret salary disclosures differently depending on whether job candidates *volunteered* unprompted, rather than provided disclosure in response to a question. We analyze this question in Appendix H and summarize here: The differences between prompted and unprompted disclosures are relatively small, and cannot be rejected from zero. We *can* reject large effects. If anything, silence in response to a prompt lowers beliefs about candidate quality more than silence without a prompt.

These results do not mean that prompts do not matter. Rather, it suggests that the prompts affect outcomes mainly through the effects on workers' choices to disclose (or not), rather than through the interpretation of disclosures (when they happen). Our supplemental survey results suggest that asking (or not) has a large effect on candidates' willingness to disclose.

Ban Simulations. We can use our experimental data to simulate the effects of banning salary history questions. To capture realistic patterns of voluntary disclosures, we combine our field experiment with our large survey of the American working public from our accompanying survey. Details of these simulations are in Appendix K. An important caveat is that these simulations do not capture how candidates and employers may respond along other margins (for example, workers changing their patterns of applications, or employers changing their search intensity).

Our simulations suggest that salary history bans reduce gender inequality in salary offers (conditional on callback). The ban increases women's salary offers from 91% of men's offers to 97%. However, they achieve this in part by reducing the salaries of all workers, and particularly men.⁵¹ Despite our earlier results, we find no detectable effects on callback inequality. One reason is that the groups whose callbacks are most impacted by disclosure have the highest rate of policy noncompliance (i.e., volunteering salary or refusing to, no

⁵¹In Table K4 we find a ban decreases men's salaries by \$8,299 while the decrease for women is statistically insignificant.

matter what the state of the ban). Finally, we study “partial” bans (used in some states, Section 2.2) that allow employers to ask for salary history after an initial offer. Our results suggest that this policy is less effective at reducing gender inequality in salaries.

9 Conclusion

What does a salary mean? In several classic models, workers are primarily differentiated by their outside options.⁵² 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.”⁵³ 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, past wages could anchor today’s salary negotiations around historical inequalities.

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. We then execute a novel, two-sided field experiment in a real world setting.

We have three main empirical results. First, employers make negative inferences about silent candidates. Candidates who do not disclose their salary are assumed to have lower-than-average quality and lower outside options, and are given lower salary offers. Disclosure choices have a large impact on recruiters’ beliefs about candidates’ outside offers, and a smaller impact on recruiter beliefs about candidates’ quality (measured by WTP) — but both are affected in the same direction. Recruiters are somewhat more lenient towards silent women.

Second, higher salaries increase recruiter’s beliefs about quality and competing offers. This leads to higher salary offers. However, we find differences in how recruiters assess higher or lower salaries. Recruiters discount male wage premiums, and place greater weight on wage residuals.

⁵²For example, models of monopsony (Manning, 2003; Card et al., 2018), efficiency wages (Lazear et al., 2016) and some models of on-the-job search (Burdett, 1978).

⁵³For models with private worker ability see Gibbons and Murphy (1992) and Oyer and Schaefer (2011), and papers about market-based tournaments (Waldman, 2013; DeVaro and Kauhanen, 2016).

Finally, salary histories can be so high that it makes a worker unattractive. In our setting, we find evidence that disclosing workers – especially those with high salaries – are *less* likely to be called back at all. This underscores our previous result about salaries signaling outside offers. At some point, outside offers can be so high that an employer could be turned off by the price. Our results highlight the potential trade-offs of disclosing and disclosing high numbers, which results in higher offers when called, but a lower likelihood of making it to the next step in the hiring process.

Although our study examines salary history, our framework provides intuition about how results could plausibly vary for other 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 competing offers. 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.
- Allon, Gad, Achal Bassamboo, and Ramandeep S Randhawa**, “Price as a Signal of Product Availability: Is it Cheap?,” *Available at SSRN 3393502*, 2012.
- Analysts, Staffing Industry**, “RPO Market Developments,” 2017.
- Arrow, Kenneth J**, “The Theory of Discrimination,” in “In Discrimination in Labor Markets” Citeseer 1973.
- Ashenfelter, Orley C, Henry Farber, and Michael R Ransom**, “Labor market monopsony,” *Journal of Labor Economics*, 2010, 28 (2), 203–210.
- Avivi, Hadar, Patrick M Kline, Evan Rose, and Christopher R Walters**, “Adaptive Correspondence Experiments,” Technical Report, National Bureau of Economic Research 2021.
- 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.
- Bertrand, Marianne and Esther Duflo**, “Field experiments on discrimination,” in “Handbook of economic field experiments,” Vol. 1, Elsevier, 2017, pp. 309–393.
- 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.
- Blackaby, David, Alison L Booth, and Jeff Frank**, “Outside offers and the gender pay gap,” *The Economic Journal*, 2005, 115 (501), F81–F107.
- Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope**, “Inaccurate Statistical Discrimination,” Technical Report, National Bureau of Economic Research 2019.
- , **Peter Hull, and Alex Imas**, “Systemic Discrimination: Theory and Measurement,” Technical Report, Working Paper 2023.
- 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.
- 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.
- Corbae, Dean and Andrew Glover**, “Employer credit checks: Poverty traps versus matching efficiency,” Technical Report, National Bureau of Economic Research 2018.

- 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**, “Salary History Bans and Unravelling,” *Working Paper*, 2022.
- **and Patryk Perkowski**, “Delegation in Hiring: Evidence from a Two-Sided Audit,” *Working Paper*, 2021.
- Crosen, Rachel and Uri Gneezy**, “Gender differences in preferences,” *Journal of Economic literature*, 2009, 47 (2), 448–74.
- Cullen, Zoë B and Bobak Pakzad-Hurson**, “Equilibrium Effects of Pay Transparency,” Technical Report, Working Paper 2021.
- Davis, Jesse, Paige Ouimet, and Xinxin Wang**, “Do Salary History Bans Limit Discrimination?,” *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.
- 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,” Technical Report, National Bureau of Economic Research 2019.
- **, Muriel Niederle, and Lise Vesterlund**, “Knowing when to ask: The cost of leaning in,” *Journal of Political Economy*, 2020, 128 (3), 816–854.
- Faberman, R Jason, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa**, “Job search behavior among the employed and non-employed,” Technical Report, NBER 2017.
- 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, Nathaniel Pattison et al.**, “Who benefits from bans on employer credit checks?,” *Dep. Work. Pap*, 2017, 1704.
- 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,” *Available at SSRN 2898827*, 2017.
- 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.
- 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.
- Kuhn, Peter J and Kailing Shen**, “What Happens When Employers Can No Longer Discriminate in Job Ads?,” Technical Report, National Bureau of Economic Research 2021.
- Kuhn, Peter, 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.

- , **Azeem M Shaikh, and Yang Xu**, “Multiple hypothesis testing in experimental economics,” *Experimental Economics*, 2019, 22 (4), 773–793.
- 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,” 2020.
- 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.
- Meli, Jeffrey and James C Spindler**, “Salary History Bans and Gender Discrimination,” Available at SSRN 3361431, 2019.
- Milgrom, Paul and John Roberts**, “Price and advertising signals of product quality,” *Journal of political economy*, 1986, 94 (4), 796–821.
- Milgrom, Paul R**, “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 1981, pp. 380–391.
- Murciano-Goroff, Raviv**, “Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers,” *Working Paper*, 2017.
- 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.
- 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.
- Phelps, Edmund S**, “The statistical theory of racism and sexism,” *The american economic review*, 1972, 62 (4), 659–661.
- Reuben, Ernesto, Paola Sapienza, and Luigi Zingales**, “How stereotypes impair women’s careers in science,” *Proceedings of the National Academy of Sciences*, 2014, 111 (12), 4403–4408.
- Robinson, Joan**, *The economics of imperfect competition*, Springer, 1933.
- Roussille, Nina**, “The central role of the ask gap in gender pay inequality,” *Working Paper*, 2020.
- Sherman, Eliot, Raina Brands, and Gillian Ku**, “The Effect of Salary History Bans on Gender-Based Disparities,” *Working Paper*, 2022.
- Sinha, Sourav**, “Salary History Ban: Gender Pay Gap and Spillover Effects,” Available at SSRN 3458194, 2020.

- Sran, Gurpal, Felix Vetter, and Matthew Walsh**, "Employer Responses to Pay History Inquiry Bans," *Working Paper*, 2020.
- 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.
- Wolinsky, Asher**, "Prices as signals of product quality," *The review of economic studies*, 1983, 50 (4), 647–658.

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 evaluations 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

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>

the recruiter. We offered to pay our subjects the hourly rate posted on their profile. We 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 application assessments and discussion 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 N. 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 key words, then we examined a random sample of approximately 5,000 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 wrote 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

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 interest in HR or hiring through the key words 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>

experience for recruiters is actually more stringent than a typical requisition for a recruiter. 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, 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

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 manually coding of this variable 92% of the time, while a recruiter's self-report of identifying as white matched our manually coding 87% of the time.

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 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 subtreatments 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 and Clustering. As described in Section 5.4, we clustered randomization. The clustered randomization produced 22 subtreatments, where a subtreatment 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

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.

E Composite Outcomes

We combine some of the outcomes generated by recruiters into composite outcomes in the following ways. These measures were assembled by us from the recruiters' assessments; the recruiters were not asked to assemble these. The recruiters' underlying assessments are listed in Section 5.2.

Employer Surplus. We measure employer surplus as WTP minus the suggested salary offer. This corresponds to $\mathbb{E}[v - s^*]$ in our conceptual framework; and it appears in Equation 3. It is also the topic of the monotonicity requirement in Section 3 (Assumption 2). The direction of this monotonicity determines whether high salary history candidates face a trade-off between receiving a callback at all (and sacrificing the salary amount) versus preferring a higher amount (and a lower chance of receiving a callback at all).

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 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.

Outside Offer Range. The outside offer range is equal to the difference between the 95th percentile and the 5th percentile of the distribution of outside offers. This gives us information about how wide-ranging the recruiter’s beliefs about outside offers are (for a given candidate). Larger ranges indicate more diffuse beliefs. This is akin to the variance of the belief distribution of η (the best outside offer).

Probability of Accepting TIOLI Offer. To estimate this probability, we fit a logistic curve through the reported 5th, 50th and 95th percentiles of the outside offer distribution.¹² We used this fitted model to predict where the salary offer made by the recruiter falls on that curve. In our conceptual framework, it is represented as $\mathbb{E}[s^* > \eta]$ in Equations 2 and 3.

Probability Accept \times Surplus. The expected surplus for the employer can be thought of as the probability a specific salary offer is accepted (described above) multiplied by the surplus should that take place. In our conceptual framework, it is represented as $\mathbb{E}[(s^* > \eta)(v - s^*)]$ in Equations 2 and 3.

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 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 primary variations of interest for the experiment.

Previous firms were chosen from the top firms that hire software engineers.¹⁵ To ascer-

¹²Similar exercises using probit curves produced results with correlations of 0.99 with the logistic approach.

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.

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>.

tain 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 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 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 is 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 postcollege job started shortly after college graduation and had a total tenure of between

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, Calif.

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.

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

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 modeled after the résumés of real software engineers.

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 Current Empl (10K)	College Grad. Year	Disclosed Salary Current Empl (10K)	Amazon	Facebook	IBM	% of Sample
All Candidates	0.50	9.97	2013.66	9.71	0.06	0.13	0.09	100.0
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
No Disclosure	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

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.

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. In specifications that do not include the amount of salary disclosed, e.g. Equation 4, to account for potential spillovers we control for the number of other applicants in the packets whose salaries *are* disclosed (this can be either 0, 3, 4, or 7). Each line of the data is a single job candidate. In specifications that do include the amount disclosed, e.g. Equation 5 we further 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

Table H1: Average Effect of Disclosing Salary for Prompted versus Unprompted Disclosure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP (10k)	WTP (10k)	Outside Offer 50th %ile	Outside Offer 50th %ile	Offer (10k)	Offer (10k)	Outside Offer Range	Outside Offer Range
Salary Disclosed	0.60** (0.21)	0.90*** (0.24)	0.82*** (0.19)	1.27*** (0.22)	0.64** (0.21)	1.02*** (0.23)	-0.45* (0.21)	-0.52* (0.22)
Salary Disclosed x Prompt	0.15 (0.29)	0.09 (0.33)	0.05 (0.27)	0.04 (0.31)	0.17 (0.28)	0.09 (0.31)	-0.88* (0.41)	-1.44+ (0.75)
Prompt on Application	-0.03 (0.26)	0.10 (0.28)	0.20 (0.25)	0.29 (0.26)	-0.01 (0.25)	0.07 (0.27)	0.66 (0.41)	1.23 (0.76)
Female x Disclosed		-0.61** (0.19)		-0.91*** (0.17)		-0.76*** (0.17)		0.15 (0.11)
Female x Salary Disclosed x Prompt		0.12 (0.27)		0.03 (0.23)		0.16 (0.25)		1.13 (0.78)
Female x Prompt on Application		-0.25 (0.21)		-0.18 (0.18)		-0.15 (0.19)		-1.14 (0.77)
Mean Unprompted Non-Disclosers:								
<i>All</i>	10.27	10.27	9.35	9.35	9.61	9.61	2.57	2.57
<i>Male</i>	10.40	10.40	9.44	9.44	9.72	9.72	2.66	2.66
<i>Female</i>	10.14	10.14	9.26	9.26	9.50	9.50	2.49	2.49
R ²	0.18	0.18	0.25	0.26	0.20	0.21	0.02	0.02
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: This Table mimics Table 2 and adds in controls and interactions with whether the disclosure was prompted or unprompted. All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in notes to Table 1. Salary Amounts and outcomes measured in dollars (e.g. WTP, offer) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

H Prompted versus Unprompted Disclosures

In this section we examine the asymmetric effects of *unprompted* disclosures within our field experiment. We measure this by estimating Equation 4 with prompt interactions, and we present results in Table H1. We also estimate Equation 5 with prompt interactions on the amounts and choice to disclose in Table H2. For most of our outcomes, we find relatively small, statistically insignificant differences. Although the differences are insignificant for most outcomes, the direction of the effects tells a common story: prompted disclosures matter more than unprompted ones.

This may have a simple explanation: prompted disclosures are more noticeable than unprompted ones. The prompt may direct visual and strategic attention to the disclosure behavior. By contrast, unprompted disclosures in our experiment require a recruiter to read the “additional information” section, notice the disclosure, and realize its significance. Although this may be an artifact of our experimental setup, similar results may happen more broadly. An unprompted disclosure to the wrong person – an interviewer instead of an HR person, or an HR person instead of the boss – may not reach the key decision maker. Additionally, the existence of a prompt for salary history on the job application may signal that an employer values this information, and thus our recruiters may rely on it more.

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.

Table H2: Average Effect of Disclosing by Salary Amount for Prompted versus Unprompted Disclosure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Offer	Outside Offer	Offer	Offer	Outside Offer	Outside Offer
	(10k)	(10k)	50th %ile	50th %ile	(10k)	(10k)	Range	Range
Salary Disclosed	-6.21***	-6.20***	-7.18***	-7.70***	-6.65***	-6.50***	0.04	-0.96
	(0.63)	(0.73)	(0.58)	(0.64)	(0.58)	(0.66)	(0.90)	(1.78)
Disclosed x Amt Disclosed (10K)	0.70***	0.71***	0.81***	0.86***	0.75***	0.75***	-0.01	0.07
	(0.06)	(0.07)	(0.06)	(0.06)	(0.06)	(0.06)	(0.11)	(0.16)
Salary Disclosed x Prompt	-0.40	-1.23+	-0.63	-0.47	-0.02	-0.90	-0.97+	-1.24
	(0.58)	(0.68)	(0.49)	(0.59)	(0.63)	(0.53)	(0.53)	(1.02)
Disclosed x Amt Disclosed (10k) x Prompt	0.06	0.07	0.07	0.02	0.02	0.05	0.01	-0.02
	(0.05)	(0.06)	(0.05)	(0.06)	(0.05)	(0.06)	(0.04)	(0.05)
Prompt on Application	-0.15	0.00	0.07	0.00	-0.14	0.00	0.65	0.00
	(0.29)	(.)	(0.27)	(.)	(0.28)	(.)	(0.43)	(.)
Female x Disclosed		0.28		0.79		-0.04		0.73
		(0.67)		(0.57)		(0.61)		(1.43)
Female x Disclosed x Amt Disclosed (10K)		-0.00		-0.06		0.02		-0.05
		(0.07)		(0.06)		(0.06)		(0.13)
Female x Salary Disclosed x Prompt		0.04		-0.80		0.31		0.70
		(0.76)		(0.63)		(0.67)		(0.91)
Female x Prompt on Application		-0.27		-0.18		-0.16		-1.14
		(0.21)		(0.17)		(0.19)		(0.78)
Female x Disclosed x Amt Disclosed (10k) x Prompt		0.03		0.10		-0.00		0.04
		(0.07)		(0.06)		(0.07)		(0.06)
Mean Unprompted Non-Disclosers:								
<i>All</i>	10.27	10.27	9.35	9.35	9.61	9.61	2.57	2.57
<i>Male</i>	10.40	10.40	9.44	9.44	9.72	9.72	2.66	2.66
<i>Female</i>	10.14	10.14	9.26	9.26	9.50	9.50	2.49	2.49
R ²	0.36	0.37	0.52	0.52	0.42	0.43	0.02	0.02
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: This Table mimics Table 4 and adds in controls and interactions with prompts. All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects. This table shows estimates from versions of Equation 5. Dependent variables are listed in the column header. Salary Amounts and outcomes measured in dollars (e.g. WTP, offer) are in \$10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.00$

We also investigate whether unprompted disclosure has gendered effects. Given gender stereotypes and cultural expectations, one may wonder whether unprompted disclosures by women might evoke negative reactions. Our estimates in Tables H1 and H2 rule out large gendered effects.

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 of the recruiters after they completed the main task, and we do find 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 concurs 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 [M1](#), [M2](#), and [M3](#) 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 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 and to be similar to candidates who disclose around the 25th percentile of workers with the same observables. This suggests our result is not likely an artifact of recruiter inexperience or lack of knowledge of market wages.

J Silence and Callbacks

Table J1: Average Effect of Disclosing for Callback Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Callback	Callback	Surplus	Surplus	p(accept)	p(accept)	p(accept) x Surplus	p(accept) x Surplus
Salary Disclosed	-0.01 (0.02)	-0.04 (0.03)	-0.05 (0.04)	-0.11* (0.04)	-0.01 (0.02)	-0.04 (0.03)	-0.01 (0.03)	-0.06+ (0.03)
Female x Disclosed		0.07+ (0.04)		0.12** (0.04)		0.06** (0.02)		0.09** (0.03)
Female Disclosure Effect:								
<i>Total</i>		0.02		0.01		0.01		0.03
<i>p-value</i>		0.48		0.80		0.57		0.30
Mean Non-Disclosers:								
<i>All</i>	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
<i>Male</i>	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
<i>Female</i>	0.62	0.62	0.60	0.60	0.56	0.56	0.32	0.32
R ²	0.03	0.03	0.01	0.01	0.03	0.04	0.02	0.03
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header and explained in notes to Table 1 (and in in Section 5.2 and Appendix E). Outcomes measured in dollars (e.g. surplus) are in \$10K increments. Using the multiple hypothesis testing correction from [List et al. \(2019\)](#), we find that these outcomes remain statistically significant after the MHT correction with the exception of $p(\text{accept})x\text{Surplus}$. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

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.

**Table J2: Average Effect of Disclosing for Callback Outcomes
Salary > \$70,000**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Callback	Callback	Surplus	Surplus	p(accept)	p(accept)	p(accept) \times Surplus	p(accept) \times Surplus
Salary Disclosed	-0.00 (0.03)	-0.03 (0.03)	-0.08* (0.04)	-0.12* (0.05)	-0.02 (0.02)	-0.04 (0.03)	-0.04 (0.03)	-0.06* (0.03)
Female \times Disclosed		0.07 (0.04)		0.09* (0.04)		0.05+ (0.02)		0.06+ (0.03)
Female Disclosure Effect:								
<i>Total</i>		0.03		-0.03		0.00		-0.01
<i>p-value</i>		0.32		0.42		0.92		0.83
Mean Non-Disclosers:								
<i>All</i>	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
<i>Male</i>	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
<i>Female</i>	0.62	0.62	0.60	0.60	0.56	0.56	0.32	0.32
R ²	0.02	0.02	0.01	0.01	0.03	0.03	0.01	0.01
Observations	1849	1849	1849	1849	1849	1849	1849	1849

Notes: All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 4 and mimics Table J1; the sample is restricted to candidates who did not disclose and those who disclosed more than \$70,000. See Section 7.4.1 and Figure 2 for explanation of the threshold. Dependent variables are listed in the column header and explained in notes to Table 1 (and in Section 5.2 and Appendix E). Outcomes measured in dollars (e.g. Surplus) are in \$10K increments. Robust standard errors are clustered at the recruiter level.

+ $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

We examine this in Table J1 using our full sample, and Table J2 using our sample that disclosed above \$70K. Our results on this question are imprecise but lean closer to the conclusion that “silence is good for callbacks” than the opposite, particularly for men and highly-paid candidates. Although our results about the callback variable itself are never statistically significant, we do find statistically significant results about employer surplus, and surplus $\times P_{Accept}$. Although silent workers are regarded as lower quality (Table 2, column 1), they are also cheaper to keep away from competing firms and opportunities (Table 2, column 3).

When we examine these results heterogeneously by gender, we find that female silence has a smaller effect on callbacks than male silence. Stated oppositely, on average women’s callbacks are punished less from disclosures than male callbacks. In fact, our gender interactions contain no statistically significant evidence of women being punished at all (on average) for disclosing. This is consistent with our earlier finding that a woman’s disclosure/silence choices contain less information about her underlying value as an employee.

K Effects of Salary History Bans

Until now, our paper has addressed how employers react to salary history disclosures. We now examine what our findings suggest about the public policies motivating our study:

salary history bans. The goal of these bans is to equalize outcomes by suppressing disclosures. However, bans cannot completely suppress disclosures. They forbid employer prompting, but not *voluntary* disclosure. Analyzing the effects of bans requires data and assumptions about worker compliance.

In this section, we combine the results of our field experiment with a survey of the working public in the United States. Our survey included approximately 1,000 individuals representative of the working American public (defined as Americans in the labor force between the ages of 22 and 55) and helps us identify which candidates are more likely to disclose unprompted (or refuse to disclose, even when asked). Here, we combine the main results of the survey with our field experiment with recruiters to study implications for salary history ban policies.

As described below, we model the ban as an increase in the cost of disclosure. As such, the prompt affects who discloses, and how disclosures are interpreted. These disclosures then affect employers' choices through the mechanisms we document in the main paper. Because we have already explored these mechanisms empirically and theoretically, this section focuses on the bottom-line effects of the ban. Our design allows us to separate the two major *designs* of the bans discussed in Section 2.2: the "full ban" (no salary questions, ever) and the "partial ban" (which allows salary history questions after the initial offer).

Theoretically, we interpret employer questions (or other prompts) as lowering m_i (the costs of disclosing); through this lens, bans raise this m for all candidates. Of course, some candidates will disclose, no matter whether the employer asks or not. These are candidates that have relatively attractive salaries, and who have low costs of disclosing. We call these candidates "always disclosers." Similarly, some workers ("never disclosers") will not disclose, even when asked. These are candidates with unattractive salaries and/or high costs. Last, "compliers" will disclose when asked, but will be silent when not asked. Compliers' salary histories are neither exceptionally attractive or not, and thus their disclosure costs are pivotal. We summarize this typology in the 2×2 matrix in Figure K1.²⁴

K.1 Estimation Strategy

Our strategy for simulating the potential impacts of a salary history ban contains several components.

K.1.1 Differentiating Prompted vs. Unprompted Disclosures We begin by measuring whether recruiters interpret salary disclosures differently depending on whether job candidates *volunteered*, rather than provided disclosure in response to a question. We analyze this question in depth in Appendix H and summarize here: The differences between prompted and unprompted disclosures are relatively small, and cannot be rejected from zero. We *can*

²⁴In principle, a fourth type exists: "Defiers," who are silent when asked and volunteer when not asked. However, our theory suggests why this group would not exist: if the ban raises costs of disclosing, all candidates should be less likely to disclose, not more. When we investigate the existence of defiers in our survey data, we find they are less than 0.5% of the U.S. workforce.

Figure K1: Ban Compliance Types

		Salary Question Banned	
		Discloses	Silent
Salary Question Asked	Discloses	Always Discloser	Ban Complier
	Silent	Ban Defier	Never Discloser

Notes: This figure labels candidates based on their disclosure choices, both in the presence and the absence of a salary history question from employers.

reject large effects. If anything, silence in response to a prompt lowers beliefs about candidate quality (compared to silence without a prompt). This small difference does not mean that prompts do not matter. Rather, it suggests that the prompts affect outcomes mainly through the effects on worker disclosure behavior, not through the interpretation of disclosures.

K.1.2 Outcomes One could combine outcomes in a number of ways to represent workers’ well-being overall. Instead, we separately present results about callbacks and salary amounts (conditional on a callback).²⁵ By presenting results on callbacks and salary amounts separately, we invite readers to import their own policy objectives about how these outcomes should be weighed. Even callbacks attached to low salaries may still be useful to workers if they grant flexibility and/or negotiating power with a current employer.²⁶

K.1.3 Differentiating Full Versus Partial Ban As discussed previously, there are two types of bans, which we call “partial” and “full” bans. Partial bans prevent employers from seeking histories only until the first offer has been made. In contrast, full bans prevent salary histories from *ever* being sought. Under both ban scenarios candidates can volunteer their salary information at any time.

For the outcomes about who is called back, outcomes for a full or partial ban are the same.

²⁵This means we restrict our sample to those who would be suggested for a callback, and look at the offers they might enjoy were they to be made a job offer.

²⁶One could instead present callbacks \times offer amount (“expected salary”). However, the outside option for a searching worker might not be zero if a job application fails. Many workers, including our fictitious candidates, might be searching for a job while currently employed and as such, lack of a callback is not a zero outcome.

Callback decisions under the full or partial bans are made only using unprompted information. However choices about salary can differ between the full and partial ban. In the partial ban, employers can ask salary history questions after making the first offer. As such, we modify our original data to simulate the effect of a partial ban. We use the callback decisions from our observations where there is no prompt on the application, and we use the offers from the observations where salary histories were asked for with a prompt. This simulates the scenario where an employer cannot ask for salary history information before the choice of whom to call back is made. But the employer *can* ask afterward about salary history and utilize any resulting information before making a final offer.

K.1.4 Regression specifications To estimate the effect of asking salary history questions with a prompt, we use Equation 8 below. Standard errors are robust and clustered at the recruiter level. The equation contains interactions that estimate the effects of these questions separately for men and women.

$$y_{i,j} = \alpha + \beta_1 \text{SalaryHistoryAsked}_j + \beta_2 \text{Female}_{i,j} + \beta_3 \text{Female}_i \times \text{SalaryHistoryAsked}_j + \epsilon_j \quad (8)$$

Notice that Equation 8 does not include a variable for whether the candidate disclosed. This is because our strategy in this section is to model candidates’ disclosure decisions—and employers’ inferences from them—as potentially downstream from the prompt. As described above, simulations of full or partial bans use the same observations if the outcome, $y_{i,j}$, is a callback, but different sets of observations when the outcome is later in the hiring process (like a TIOLI amount for those called back).

K.1.5 Regression weights Our experiment measures recruiters’ reactions to all sets of compliance behaviors for each candidate, both with and without the prompt. 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. To incorporate assumptions about worker compliance, we estimate Equation 8 using regression weights. These weights place more (or less) weight on observations resembling expected candidate behavior.

To show an example of this use of weights, suppose we wanted to assess the ban assuming that all subjects were “compliers.” We would assign a weight = 1 for all observations in which the candidate is prompted and discloses, and (similarly) a weight = 1 for observations in which the candidate is *not* prompted and does not disclose. All other observations would receive a weight of zero.²⁷ In this example of compliers, the coefficient on the prompt captures the effect of all workers changing from non-disclosers to disclosures when prompted. In footnote 28, we review a weighting scheme as if everyone were an “always-discloser.”²⁸

²⁷This effectively drops all observations in which a candidate is asked and stays silent, or volunteers without asking.

²⁸Suppose we wanted to assess the ban assuming that all subjects were “always-disclosers.” We would receive a weight= 1 for all observations in which the candidate is prompted and discloses, and (similarly) a weight= 1 for observations in which the candidate is *not* prompted and yet discloses (volunteers unprompted). All other observations would receive a weight of zero. This effectively drops all observations in which a candidate does not disclose, leaving only observations where the candidate discloses (either with prompting, or not).

We aim to employ weights for Equation 8 reflecting the true probabilities of the disclosure behavior for each of our candidates. If women like Jessica from Oracle (CandidateId #5 in our experiment) tend to be compliers, we would place a greater weight on her observations that include disclosure (when prompted) and non-disclosure (when not prompted). Below, we discuss our strategy for obtaining the weights. Using these weights, we can use Equation 8 to estimate how the ban will affect aggregate outcomes of our experiment (for example, the overall level of men’s and women’s salary offers, and their ratio).

K.2 Surveys of Workers and Regression Weights

To estimate each of our candidates’ likely compliance behavior, we conducted a survey of over 1,000 American workers ages 22 to 55 using a survey company.²⁹ The surveys asked respondents about their demographics, and whether they would disclose their salary when asked (or volunteer when not asked).³⁰ Using this data, we can identify each survey respondent’s disclosure type (“always-discloser,” “never-discloser,” etc.) by asking about candidate disclosure behavior when prompted by an employer (and not). We also asked how workers in the survey volunteer (or not) in scenarios where asking is illegal.³¹ Our survey also asked subjects to identify their gender, their overall income, their industry and occupation, and whether they were relatively well paid (or not) compared to other people in the same job at the same company.

These covariates helped us link survey respondents to fictitious candidates in our experiment with similar characteristics. Using these data points, we developed a map between our survey responses and the characteristics of our job candidates. Table K1 displays some descriptive regressions of these mappings. Always-disclosers make up 28% of our survey sample. They are more likely to be male and are slightly higher-paid.³² A majority, 52%, of our survey sample are compliers. Compliers have the opposite set of correlations: more female, and in lower paying occupations, lower paid conditional on jobs and industries, and less likely to be paid more than peers. Never disclosers make up 19% of our sample; we find they are more likely to report high salaries within their firm, but to work in lower paying

²⁹We used the company Prolific Academic, <https://www.prolific.co/>. In a separate paper (Cowgill et al., 2022), we document this survey instrument and analyze the data included in this paper in more detail.

³⁰The exact wording of the latter question was, “Imagine that no one involved in the hiring process has asked you about your most recent salary. However, you can legally disclose this information voluntarily. Would you tell them your most recent salary?” For the first question, it was, “Imagine it is perfectly legal for someone involved in the hiring process to ask your most recent salary. If someone asks, would you tell them your most recent salary?” For a subset of respondents we have information about their choices in two real job searches in which they were asked and not asked for salary history. In addition to these hypothetical scenarios, we also asked about the subjects’ actual disclosures when they encountered questions in their job searches. Answers to the hypothetical questions and real questions are positively correlated. All these results are available from authors upon request.

³¹When salary questions have been banned, workers may be more (or less) willing to volunteer – compared to settings where they are allowed, but companies decline to ask. Our surveys explicitly address this possibility. We find that when an employer does not ask about 70-74% of workers do not disclose regardless of three scenarios 1) whether asking is banned by law, 2) is legal but the employer chooses not to, or 3) if the legality is ambiguous (Cowgill et al., 2022).

³²They work in slightly higher paying jobs and industries.

Table K1: Who are Complier Types?

Panel A: Always Disclosers

	Always Discloser	Always Discloser	Always Discloser	Always Discloser	Always Discloser	Always Discloser	Always Discloser	Always Discloser
Female	-.12*** (.028)					-.11*** (.028)	-.11*** (.03)	-.13*** (.045)
High Salary w/in Firm		-.017 (.028)				-.016 (.028)	-.012 (.028)	-.035 (.043)
Salary (Normalized)			.054*** (.016)			.048*** (.012)	.046*** (.011)	.067*** (.023)
Occupation's Average Salary (Norm)				.031** (.015)				
Industry's Average Salary (Norm)					.021 (.015)			
Female × High Salary w/in Firm								.038 (.059)
Female × Salary (Norm)								.053 (.087)
High Salary w/in Firm × Salary (Norm)								-.0003 (.068)
Fem. × High Salary @Firm × Salary (Norm)								-.086 (.11)
Industry FEs							Y	Y
Occupation FEs							Y	Y
R ²	.018	.00038	.013	.0049	.0022	.029	.059	.061
Observations	1,006	1,006	1,006	1,006	1,006	1,006	1,005	1,005

Panel B: Ban Compliers

	Ban Complier	Ban Complier	Ban Complier	Ban Complier	Ban Complier	Ban Complier	Ban Complier	Ban Complier
Female	.11*** (.031)					.11*** (.031)	.12*** (.034)	.11** (.05)
High Salary w/in Firm		-.034 (.032)				-.035 (.031)	-.043 (.032)	-.041 (.046)
Salary (Normalized)			-.045*** (.016)			-.04*** (.013)	-.044*** (.015)	-.076*** (.028)
Occupation's Average Salary (Norm)				-.0074 (.016)				
Industry's Average Salary (Norm)					.00013 (.016)			
Female × High Salary w/in Firm								.008 (.066)
Female × Salary (Norm)								-.079 (.1)
High Salary w/in Firm × Salary (Norm)								.019 (.071)
Fem. × High Salary @Firm × Salary (Norm)								.11 (.12)
Industry FEs							Y	Y
Occupation FEs							Y	Y
R ²	.012	.0011	.0076	.00022	7.1e-08	.019	.047	.05
Observations	1,006	1,006	1,006	1,006	1,006	1,006	1,005	1,005

Panel C: Never Disclosers

	Never Discloser	Never Discloser	Never Discloser	Never Discloser	Never Discloser	Never Discloser	Never Discloser	Never Discloser
Female	.014 (.025)					.012 (.025)	-.0018 (.027)	.023 (.038)
High Salary w/in Firm		.048* (.025)				.048* (.025)	.051** (.025)	.07** (.035)
Salary (Normalized)			-.0088 (.0066)			-.0082 (.0066)	-.0026 (.0074)	.0093 (.02)
Occupation's Average Salary (Norm)				-.021* (.012)				
Industry's Average Salary (Norm)					-.018 (.012)			
Female × High Salary w/in Firm								-.042 (.051)
Female × Salary (Norm)								.023 (.081)
High Salary w/in Firm × Salary (Norm)								-.031 (.06)
Fem. × High Salary @Firm × Salary (Norm)								-.0068 (.1)
Industry FEs							Y	Y
Occupation FEs							Y	Y
R ²	.00033	.0037	.00053	.0028	.0022	.0045	.024	.025
Observations	1,006	1,006	1,006	1,006	1,006	1,006	1,005	1,005

Notes: + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

occupations.

Using this data, we construct eight cells ($\{\text{male, female}\} \times \{\text{High Wage Firm, Low}\} \times \{\text{High Salary within Firm, Low}\}$). The distribution of disclosure types of all eight groups can be browsed in Table K2. The proportion of the population in each of these cells is relatively uniform between 11.1% and 15.4%. Across these cells, women are more likely to be compliers, and are less likely to be always-disclosers. Respondents working in high wage firms are more likely to be always-disclosers and less likely to be compliers. Within each cell, never-disclosers are always the smallest proportion of a compliance type.

Each of the candidates in our field experiment can be mapped back to one of these cells. We thus merge the contents of Table K2 with our field experimental data, giving weights to the experimental observations that the survey indicates are more likely. The final columns of Table K2 summarize how much each group gains/loses from disclosing (on average).

Table K2: Complier Type Cells

Gender	Type of Cells		Proportion of Cell that is an ...				Mean Disclosure Δ	
	High/Low Wage Firm	High/Low Salary vs Peers @Firm	Always Discloser	Never Discloser	Ban Complier	% of Sample	Callback Prob	TIOLI Offer
Female	High	High	0.28	0.24	0.48	11.6	+0.02	+\$21.7K***
Female	High	Low	0.23	0.21	0.55	11.1	+0.06	+\$2.7K
Female	Low	High	0.18	0.19	0.63	15.4	+0.03	+\$834
Female	Low	Low	0.20	0.15	0.65	11.5	-0.04	-\$13.3K***
Male	High	High	0.33	0.23	0.43	15.0	-0.15**	+\$28.3K***
Male	High	Low	0.42	0.15	0.43	11.4	-0.003	+\$12.7K***
Male	Low	High	0.31	0.19	0.49	11.6	+0.05	+\$7.38K***
Male	Low	Low	0.31	0.15	0.54	12.2	-0.05	-\$9.79K***

Notes: This table aggregates data from our survey of 1,006 US job seekers (middle columns) and our experiment (final two columns). It reports the distribution of compliance types by gender, whether a person works at a high wage firm, and whether they self-report having a salary above the median within their firm. The final columns report the average changes in outcomes from candidate disclosures in our experiment. An “Always Discloser” always reports their salary whether prompted or not, a “Never Discloser” never reports, and a “Ban Complier” will disclose if prompted.

K.3 Ban Simulation Results

K.3.1 Ban Effects on Callbacks Table K3 shows the effects of banning employers from prompting salary history disclosure on whether women and men in our field experiment are selected to receive a callback. Under either a partial or a full ban, the recruiter would only observe the salary information on the job application without a prompt at the time that callback choices are made, and as such the results are the same under either type of ban. Either with or without a ban 64% of women and 63% of men are recommended for a callback. There is no gender gap in callbacks before the ban, and there continues to be no gender gap after the ban in our setting with a preexisting gender wage gap.

These results are similar to a number of studies which observe small or nonexistent changes

Table K3: Effect of Bans (Full or Partial) on Callbacks

Callbacks

	Women	Men	Ratio
No Ban	.64 (.023)	.63 (.023)	1 (.045)
Ban	.64 (.029)	.63 (.024)	1 (.05)
Ban-No Ban	-.0034 (.037)	-.0027 (.033)	-.00096 (.067)
<i>p</i> -value	.93	.93	.99

Notes: This table shows the effect of a salary history ban on whether a candidate was recommended for a callback. Under both a full or partial ban the callback decision is always made when only the information from the application is available either with a prompt under No Ban, or without a prompt under Ban. Standard errors are robust.

in employment or job changing from bans (Bessen et al., 2020; Sinha, 2020; Hansen and Mc-Nichols, 2020; Mask, 2020).³³ However, these results seem to contrast with our earlier results in which silence *increased* callbacks (Section 7.4). Table K2 reconciles these findings: the groups whose callbacks are most impacted by disclosing have the highest rate of noncompliance (voluntary disclosure and refusal).

Specifically, well-paid men at high-wage firms are 15 percentage points less likely to receive a callback when they disclose. The ban should help their callback rates by silencing them. However, this group is full of always disclosers (33%) and never disclosers (23%). Only 43% answer only when asked—the lowest percentage of compliers of any group.

K.3.2 Effects of a Ban on Salary Offers Conditional on Callback Table K4 presents the effects of full bans, where salary disclosure may never be prompted, and partial bans, where an inquiry may be made later in the hiring process. The story is similar for both types of bans: Bans close the gender gap. In the left-hand panel of Table K4, we see that the ratio of annual salary conditional on callback for women to men is 0.91 before a ban, and rises to 0.97 after a ban ($p = 0.00$), meaning that women and men are almost equal after the ban is in place. A partial ban has the same pattern of effects, but the magnitude is significantly smaller and less precise. The ratio of female to male is 0.91 before a ban, and rises to 0.92 after a ban ($p = 0.13$).³⁴

However, *how* bans close the gender gap is important. Bans *lower* salary offers for both women and men. But a full ban harms men by an average of \$8,299 ($p = 0.00$) while women

³³Some other studies find changes in job transitions either overall (Sran et al., 2020) or for those who entered the labor force during a recession (Mask, 2020).

³⁴The ratio of female to male salaries disclosed is 0.85 in our experiment by construction. Under all the policy regimes studied, the recruiters in our experiment narrowed the wage gap.

Table K4: Full vs. Partial Ban: Effect on Annual Salaries, Conditional on Callbacks

Salary | Callbacks (Full Ban)

	Women	Men	Ratio
No Ban	101824.98 (1111.61)	112220.58 (1315.39)	0.91 (0.01)
Ban	100377.55 (1576.33)	103921.13 (1349.36)	0.97 (0.01)
Ban-No Ban	-1447.43 (1928.86)	-8299.45 (1884.42)	0.06 (0.01)
<i>p</i> -value	0.45	0.00	0.00

Salary | Callbacks (Partial Ban)

	Women	Men	Ratio
No Ban	101824.98 (1111.61)	112220.58 (1315.39)	0.91 (0.01)
Ban	97966.75 (319.37)	106061.10 (478.83)	0.92 (0.01)
Ban-No Ban	-3858.23 (1156.58)	-6159.48 (1399.83)	0.02 (0.01)
<i>p</i> -value	0.00	0.00	0.13

Notes: This table shows the effect of a salary history ban on the salary offer when a candidate was recommended for a callback. The left panel shows this for a full ban and the right panel is for a partial ban. A “Full Ban” means a ban where salary history may not be asked at any stage in the hiring process. A “Partial Ban” means a ban of prompting job candidates to disclose their salary history on the job application, but being able to seek salary information at a later stage in the hiring process. Standard errors are robust.

only lose \$1,447 ($p = 0.45$). The results are even more harmful to women for a partial ban, with women losing \$3,858 ($p = 0.00$) and men losing \$6,159 ($p = 0.00$) on average. This same pattern can be seen for alternative outcomes like salary \times callback (Table K5). For either type of ban, the gender gap is closed by greater harm to men than women. Policymakers may have hoped that salary history bans would raise salary offers for women, but our field experiment shows women receive lower salary offers and men receive much lower offers.

In our section on the effects of disclosure (Section 7), we found that silence lowered employers’ beliefs about candidate quality and (especially) competing offers. This resulted in lower salary offers. It is possible that a salary history ban would have no effect on the levels of silence vs disclosing (i.e., if the world was made up of only always-disclosers and never-disclosers). However, our survey shows that about 52% of the U.S. workforce are ban compliers, and as such bans increase silence. Bans lower inequality in salary offers, but do so by increasing silence. The resulting silence harms men more than women. In short, bans divide the salary offers pie more equally between male and female job candidates, but they also shrink the total size of the pie.

Table K5: Full vs. Partial Ban: Effect on Expected Salary Offers (Salary Offer \times Callback)

Salary \times Callbacks (Full Ban)

	Women	Men	Ratio
No Ban	64919.37 (2391.37)	70646.28 (2717.71)	0.92 (0.04)
Ban	63937.62 (2886.18)	64882.21 (2487.74)	0.99 (0.05)
Ban-No Ban	-981.75 (3748.16)	-5764.07 (3684.40)	0.07 (0.07)
<i>p</i> -value	0.79	0.12	0.31

Salary \times Callbacks (Partial Ban)

	Women	Men	Ratio
No Ban	64919.37 (2391.37)	70646.28 (2717.71)	0.92 (0.04)
Ban	62402.00 (2734.08)	66218.28 (2431.93)	0.94 (0.05)
Ban-No Ban	-2517.36 (3632.33)	-4428.00 (3646.95)	0.02 (0.06)
<i>p</i> -value	0.49	0.23	0.71

Notes: This table shows the effect of a salary history ban on the salary offer multiplied by a binary variable for if the candidate was recommended for a callback. The left panel shows this for a full ban and the right panel is for a partial ban. These are the effects for a “Full Ban” meaning a ban where salary history may not be asked at any stage in the hiring process. The effects for a “Partial Ban” meaning a ban of prompting job candidates to disclose on the job application, but being able to seek salary information at a later stage in the hiring process. Standard errors are robust.

L External Validity: SANS conditions

In this section we discuss 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 4 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 14% of subjects did not complete the task after being sent materials. Dropout from the study was correlated with our randomly assigned treatment arms: Recruiters were 12 percentage points less likely to complete the task if they were randomly assigned a packet that did *not* ask for salary history. In addition, recruiters were 15 percentage points less likely to complete the job if they were assigned a packet where no single candidate provided their salary history (i.e., all candidates were “silent”) even if they were asked. Although these effects are moderate in size, they are statistically significant (Table L1).

Table L1: Attrition

	Remained	Remained	Remained
Salary Prompt Appears On Job Application Form	.21*** (.029)		.12*** (.021)
At Least One Disclosure		.22*** (.03)	.15*** (.024)
Observations	298	298	298
R ²	.078	.088	.11
Mean Dep. Var	.86	.86	.86

Notes: This table studies attrition of invited subjects into the study. Of 298 subjects sent material, about 14% dropped out. In the regressions above, we study which invited subjects remained in the study (versus dropping out) as a function of their treatment assignment.

Our treatments were not designed to study recruiter retention. To be clear, this idea does not appear in any form within the conceptual framework guiding our thinking (Section 3). We discovered our attrition results only by checking standard attrition diagnostics. However, it is not hard to imagine the behavioral forces might lead to these results. Proposing wages requires more work for the recruiter without the guidance of salary histories. If candidates are not providing benchmarks, and the employer isn't even seeking them – then the task requires more effort. Some recruiters might not find the job worth undertaking, even for payment.

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, 2020). 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 unravelling would proceed, and our empirical results contain some evidence of unravelling dynamics (i.e., silence assumed to be a negative signal).

Outside of our experiment, we find additional suggestive evidence of unravelling. A report by the New York Times about salary history bans in 2021 says, “said some people [...] volunteered their salary history.” One candidate told the Times. “I prefer to be direct about what I’m making.”³⁵ 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). Unravelling and other adjustments to bans raise the potential for a “voltage drop” after scaling up this policy.

³⁵<https://www.nytimes.com/2021/12/30/business/salary-negotiation-pay.html>

M Additional Empirical Analysis

Table M1: Average Effect of Disclosing for Packets with Half of Salaries Disclosed

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	WTP	WTP	Outside Offer 50th %ile	Outside Offer 50th %ile	Offer	Offer	Outside Offer Range	Outside Offer Range
Salary Disclosed	0.50*** (0.13)	0.60** (0.20)	0.61*** (0.12)	0.98*** (0.17)	0.57*** (0.11)	0.75*** (0.17)	-0.45 (0.33)	-0.87 (0.63)
Female x Disclosed		-0.21 (0.26)		-0.74*** (0.21)		-0.36 (0.23)		0.83 (0.63)
Female Disclosure Effect:								
<i>Total</i>		0.40		0.24		0.39		-0.04
<i>p-value</i>		0.01		0.11		0.01		0.70
Mean Non-Disclosers:								
<i>All</i>	10.30	10.30	9.48	9.48	9.65	9.65	2.85	2.85
<i>Male</i>	10.49	10.49	9.61	9.61	9.80	9.80	3.20	3.20
<i>Female</i>	10.11	10.11	9.35	9.35	9.51	9.51	2.50	2.50
R ²	0.08	0.07	0.11	0.12	0.11	0.11	0.00	0.00
Observations	768	768	768	768	768	768	768	768

Notes: This table shows estimates from versions of Equation 4 and mimics Table 2; the sample is restricted to data from recruiters who evaluated packets with exactly half of salaries disclosed (4 disclosed salaries, 4 non-disclosed salaries). All models include recruiter controls and candidate fixed effects. Dependent variables are listed in the column header and explained in notes to Table 1. Outcomes measured in dollars (e.g. WTP, Offer) are in \$10K increments. This table is the subset of the data presented in the main text. This table only shows data where a recruiter saw exactly half of the candidates disclosing their salary history versus all or none of the candidates disclosing as presented in Table 2. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$.

Table M2: Average Effect of Disclosing a High versus Low Salary for Packets with Half of Salaries Disclosed

	(1) WTP	(2) Outside Offer 50th %ile	(3) Offer	(4) Outside Offer Range
Disclosed 25th %ile Salary	-0.09 (0.15)	-0.04 (0.14)	-0.05 (0.13)	-0.65+ (0.38)
Disclosed 75th %ile Salary	1.13*** (0.17)	1.28*** (0.15)	1.21*** (0.15)	-0.24 (0.31)
Mean Non-Disclosers	10.34	9.64	9.72	2.65
R ²	0.22	0.30	0.25	-0.01
Observations	768	768	768	768

Notes: This table mimics Table 3 but the sample is restricted to data from recruiters who evaluated packets with exactly half of salaries disclosed (4 disclosed salaries, 4 non-disclosed salaries). All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects as described in the text. Dependent variables are listed in the column header and explained in notes to Table 1. 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. The omitted category is candidates who did not disclose a salary. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table M3: Average Effect of Disclosing by Salary Amount for Packets with Half of Salaries Disclosed

	(1) WTP	(2) WTP	(3) Outside Offer 50th %ile	(4) Outside Offer 50th %ile	(5) Offer	(6) Offer	(7) Outside Offer Range	(8) Outside Offer Range
Salary Disclosed	-5.33*** (0.74)	-5.69*** (0.88)	-5.78*** (0.71)	-5.80*** (0.82)	-5.54*** (0.66)	-5.66*** (0.78)	0.32 (1.81)	0.80 (2.79)
Female x Disclosed		0.04 (0.89)		-0.15 (0.83)		-0.27 (0.78)		-1.44 (2.37)
Female Disclosure Slope:								
Total		0.67		0.69		0.70		0.06
p-value		0.00		0.00		0.00		0.54
Mean Non-Disclosers:								
All	10.30	10.30	9.48	9.48	9.65	9.65	2.85	2.85
Male	10.49	10.49	9.61	9.61	9.80	9.80	3.20	3.20
Female	10.11	10.11	9.35	9.35	9.51	9.51	2.50	2.50
R ²	0.27	0.27	0.36	0.36	0.31	0.32	-0.01	-0.01
Observations	768	768	768	768	768	768	768	768

Notes: This table mimics Table 4 but the sample is restricted to data from recruiters who evaluated packets with exactly half of salaries disclosed (4 disclosed salaries, 4 non-disclosed salaries). All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects as described in the text. This table shows estimates from versions of Equation 4 that include interactions with gender. Dependent variables are listed in the column header and explained in notes to Table 1. Outcomes measured in dollars (e.g. WTP, Offer) 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 M4: Average Effect of Disclosing Salary for Additional Outcomes

	(1) Outside Offer 5th %ile	(2) Outside Offer 5th %ile	(3) Outside Offer 95th %ile	(4) Outside Offer 95th %ile	(5) ≥ 2 Other Offers	(6) ≥ 2 Other Offers
Salary Disclosed	1.17*** (0.14)	1.67*** (0.16)	0.25 (0.27)	0.42 (0.43)	-0.01 (0.03)	-0.03 (0.04)
Female x Disclosed		-1.01*** (0.11)		-0.34 (0.37)		0.03 (0.04)
Female Disclosure Effect:						
<i>Total</i>		0.67		0.08		0.01
<i>p-value</i>		0.00		0.65		0.82
Mean Non-Disclosers:						
<i>All</i>	8.16	8.16	11.01	11.01	0.54	0.54
<i>Male</i>	8.23	8.23	11.43	11.43	0.57	0.57
<i>Female</i>	8.09	8.09	10.59	10.59	0.50	0.50
R ²	0.27	0.28	0.06	0.06	0.03	0.03
Observations	2048	2048	2048	2048	2048	2048
	(7) Candidate Searches	(8) Candidate Searches	(9) Firm Searches	(10) Firm Searches	(11) Both Search	(12) Both Search
Salary Disclosed	-0.02 (0.03)	0.00 (0.03)	0.02 (0.03)	-0.00 (0.03)	-0.01 (0.03)	-0.04 (0.03)
Female x Disclosed		-0.04 (0.04)		0.04 (0.04)		0.07+ (0.04)
Female Disclosure Effect:						
<i>Total</i>		-0.04		0.04		0.03
<i>p-value</i>		0.23		0.23		0.41
Mean Non-Disclosers:						
<i>All</i>	0.42	0.42	0.58	0.58	0.41	0.41
<i>Male</i>	0.39	0.39	0.61	0.61	0.45	0.45
<i>Female</i>	0.46	0.46	0.54	0.54	0.37	0.37
R ²	0.03	0.03	0.03	0.03	0.00	0.00
Observations	2048	2048	2048	2048	2048	2048

Notes: This table mimics Table 2 but for additional outcomes we collected. All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 4. Dependent variables are listed in the column header. Columns 5-12 are binary outcomes: ≥ 2 Other Offers means recruiter thinks the candidate will have 2 or more outside offers (as opposed to 1 or fewer); Candidate Searches means recruiter thinks outside offer likely comes from candidate aggressively pursuing outside options; Firm Searches means recruiter thinks outside offers likely come from other firms aggressively pursuing this candidate; and Both Search means recruiter believes outside offers come from both candidate and other firms pursuit. Outcomes measured in dollars (e.g. the outside offers) 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 M5: Average Effect of Disclosing by Salary Amount for Additional Outcomes

	(1) Outside Offer 5th %ile	(2) Outside Offer 5th %ile	(3) Outside Offer 95th %ile	(4) Outside Offer 95th %ile	(5) ≥ 2 Other Offers	(6) ≥ 2 Other Offers
Salary Disclosed	-6.23*** (0.54)	-6.24*** (0.60)	-6.42*** (0.98)	-7.37*** (1.74)	-0.02 (0.15)	-0.08 (0.18)
Female x Disclosed		-0.11 (0.37)		0.91 (1.57)		0.05 (0.16)
Disclosed x Amt Disclosed	0.71*** (0.06)	0.71*** (0.06)	0.69*** (0.12)	0.74*** (0.19)	0.01 (0.01)	0.01 (0.02)
Female x Disclosed x Amt Disclosed		0.03 (0.04)		-0.01 (0.15)		0.00 (0.02)
Female Disclosure Effect:						
<i>Total</i>		0.74		0.73		0.01
<i>p-value</i>		0.00		0.00		0.40
Mean Non-Disclosers:						
<i>All</i>	8.16	8.16	11.01	11.01	0.54	0.54
<i>Male</i>	8.23	8.23	11.43	11.43	0.57	0.57
<i>Female</i>	8.09	8.09	10.59	10.59	0.50	0.50
R ²	0.54	0.54	0.12	0.12	0.05	0.05
Observations	2048	2048	2048	2048	2048	2048
	(7) Candidate Searches	(8) Candidate Searches	(9) Firm Searches	(10) Firm Searches	(11) Both Search	(12) Both Search
Salary Disclosed	0.05 (0.14)	0.10 (0.17)	-0.05 (0.14)	-0.10 (0.17)	-0.15 (0.15)	-0.12 (0.18)
Female x Disclosed		-0.01 (0.17)		0.01 (0.17)		-0.13 (0.17)
Disclosed x Amt Disclosed	-0.01 (0.01)	-0.01 (0.02)	0.01 (0.01)	0.01 (0.02)	0.02 (0.02)	0.01 (0.02)
Female x Disclosed x Amt Disclosed		-0.01 (0.02)		0.01 (0.02)		0.02 (0.02)
Female Amount Disclosed Slope:						
<i>Total</i>		-0.02		0.02		0.03
<i>p-value</i>		0.26		0.26		0.04
Mean Non-Disclosers:						
<i>All</i>	0.42	0.42	0.58	0.58	0.41	0.41
<i>Male</i>	0.39	0.39	0.61	0.61	0.45	0.45
<i>Female</i>	0.46	0.46	0.54	0.54	0.37	0.37
R ²	0.06	0.06	0.06	0.06	0.03	0.03
Observations	2048	2048	2048	2048	2048	2048

Notes: This table mimics Table 4 but for additional outcomes we collected. All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects as described in the text. This table shows estimates from versions of Equation 4 that include interactions with gender. Dependent variables are listed in the column header and explained in notes to Table M4. Salary Amounts and outcomes measured in dollars (e.g. the outside offers) 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 M6: Effect of An Extra Dollar Decomposed for Additional Outcomes

	(1) Outside Offer 5th %ile	(2) Outside Offer 95th %ile	(3) Got Only Offer	(4) ≥ 2 Other Offers	(5) Candidate Searches	(6) Firm Searches	(7) Both Search
+\$10k from Firm	0.73*** (0.06)	0.72*** (0.21)	-0.01 (0.01)	0.02 (0.02)	-0.01 (0.02)	0.01 (0.02)	0.01 (0.02)
+\$10k from Male	0.62*** (0.08)	0.23 (0.27)	-0.02 (0.02)	-0.02 (0.03)	0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)
+\$10k within Firm	0.71*** (0.06)	0.75*** (0.06)	0.01 (0.01)	0.01 (0.02)	-0.02 (0.02)	0.02 (0.02)	0.04* (0.02)
<i>p</i> F-M	0.04	0.04	0.50	0.17	0.09	0.09	0.14
<i>p</i> F-W	0.63	0.87	0.14	0.89	0.82	0.82	0.14
<i>p</i> M-W	0.14	0.04	0.13	0.25	0.07	0.07	0.02
R ²	0.54	0.12	0.01	0.05	0.06	0.06	0.03
Observations	2048	2048	2048	2048	2048	2048	2048

Notes: This table mimics Table 5 but for additional outcomes we collected. All models include recruiter and spillover controls and both candidate and sub-treatment fixed effects as described in the text. This table shows estimates from Equation 5 which decomposes additional dollars of salary disclosure into a firm-specific offset for the candidate’s employer (“+\$10k from Firm,” some firms pay higher or lower to everyone on average); a gender offset (“+\$10k from Male”, which mimics real-world gender gaps); and from having a higher or lower salary within the current firm’s distribution (“+\$10k within Firm”, note this also is combined with some random noise that was included in the salaries.). Dependent variables are listed in the column header and explained in notes to Table M4. Outcomes measured in dollars (e.g. the outside offers) are in \$10K increments. *p*-values for comparisons of coefficients within the same model are provided in the 2nd panel, where for example *p* F-M is the *p*-value testing that the coefficient from “+\$1 from Firm” = the coefficient on “+\$1 from Male”. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

N Experimental Materials

N.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.

N.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]-[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.

N.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.

N.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]-[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).

N.5 Recruiter Instructions

Instructions

Thank you for your help screening our candidates. Please read these instructions carefully and completely before you begin this task.

1 About our Hiring Needs

We are interested in finding candidates for a full-stack software engineering position at a mid-sized software start-up company. Qualified candidates should have a working understanding of hardware systems infrastructure, creating and manipulating databases, writing back-end code in one or more languages (e.g., Ruby, Java, Python, C#), and writing front-end code in one or more languages (e.g., HTML, Javascript). Other responsibilities may include project management and technical documentation. Our company has locations in several cities throughout the United States and many of our software engineers work remotely; location will be determined in consultation with the candidate after an offer has been made.

Additional details about our opening are available in section 5.B.

2 Your Task

We will provide you with candidates' responses to our online job application form. We ask that you review this information and answer a few questions. In particular, we will ask you about:

- Whether we should interview the candidate
- What salary we should offer or accept if they pass our interview
- Additional questions about potential salary ranges

At this stage, we are interested in identifying worthy candidates. In that sense, we do not have a fixed number of positions so you should let us know about any candidate in our applicant pool that would be a good match for this position.

Software engineers currently at our firm make between \$70,000 and \$120,000. You should not feel constrained by our current range, and we welcome your own research about what candidates should be paid. We also offer benefits including health insurance, stock and a performance-based annual bonus. However, our questions for you today will be about the cash component (annual base salary) of compensation only.

2 Compensation for you, our recruiter, for this task

For your assistance with this task, you will be paid hourly (with a maximum of 2 hours allowable), plus a bonus. You can read the details about the bonus calculation in the appendix to these instructions, but we'll summarize it here:

1. We care about spending recruiting energy on candidates we're likely to hire -- candidates who will impress us in interviews and will accept our offers.

2. We care about the difference between what we pay candidates and the value they bring to our company. It's worth paying more for salaries, but only if they bring more value (and/or if they're more likely to accept). We want your decisions to consider value, cost, and probability of acceptance.

We will interview all candidates you suggest. We may also interview candidates you did not suggest upon recommendation from others at our company.

Please note, we do not negotiate salaries with candidates.

3 Your Feedback about the Candidates

We will provide an online form for you to fill in your evaluations to make it easier to work together without too much back-and-forth. There will be six sets of questions about the candidates themselves, and a few quick questions about yourself.

4 Additional Information

Ultimately our staff are very busy and not available to answer questions as you review these applications.

Please do not contact any of these candidates. We are asking you only to evaluate them and send us your private assessments. Someone from our staff will take the next step with the candidates. To prohibit you from contacting them, we have blacked out their contact information in the attached application forms.

Our hiring philosophy is to make interview decisions based on what is submitted. Therefore, please do not consult any information on individual candidates outside of the packets we send you. For example, do not look up the candidates on Google or LinkedIn.

5 Appendix

The remainder of this document includes:

- Some additional details about your bonus payment.
- Additional information about the job requirements for full-stack software engineer.

5.A Exact Formula for Calculating Your Bonus

We will calculate a bonus associated with each candidate you review according to the guidelines below, and then sum them up across all candidates and pay you the full sum in addition to your hourly rate. The bonuses will be paid after we have completed our interview and hiring decisions - approximately 45 days (or sooner) after you complete the task.

For candidates who are hired, we will examine their performance and trajectory about four weeks after the candidate starts work. We'll rate the newly hired candidate on three dimensions using the one-through-three scale outlined on the next page.

We will add up the candidate's three scores, for a total score ranging between 3 and 9. We then multiply that total score by five, and subtract [the candidate's salary / 100,000]. This is your bonus for each newly hired candidate.

$$\text{Hired Candidate Bonus} = 5 \times (\text{Technical Score} + \text{Innovation Score} + \text{Leadership Score}) - \text{Salary}/100,000$$

As you know, you'll help set our workers' salaries through your feedback in this task.¹ This bonus gives you the incentive to find candidates who deliver a lot of value to our company above the salary we need to pay them.

For candidates who are NOT hired -- either because we don't make them an offer, or because they reject our offer -- your Hired Candidate Bonus for that candidate will be zero.

We will also subtract \$5 from your overall bonus for everyone you suggest interviewing who isn't hired. This is to encourage you to be a little bit selective about forwarding candidates who have a realistic shot at joining our company. If we hire someone who you didn't suggest interviewing, we'll calculate the Hired Candidate Bonus as if you suggested interviewing that candidate. Also: If you suggest interviewing a candidate and the candidate declines to be interviewed, we would count this as a failed interview.

Please note: It would (in theory) be possible to earn a negative overall bonus. If this happens, we will set the overall bonus to \$0.

¹ In one of our questions for you, we'll ask you what we should offer the candidate as a take-it-or-leave-it offer. For candidates who accept, we'll use that salary in the bonus calculation. We'll also ask you what to do if a candidate instead approaches us with a take-it-or-leave-it offer. If you guide us to accept those offers in some circumstances, then we'll use those salaries in the formula above.

Evaluation Dimensions

- A. Technical Score
- B. Innovation Score
- C. Leadership Score

Examples of Performance in Each Dimension

A. Technical Score:

Rating 1 (Low): Gaining command of all core technologies and practices used in our firm's engineering team. Able to begin developing and productionizing low to moderate complexity modules.

Rating 2 (Middle): Reasonable command of core engineering systems. Shows comfort with owning reasonably high complexity modules.

Rating 3 (High): Responsible for driving, technically designing, implementing and productionizing high impact projects with the help of teams if needed. Can own and deliver on very large mission-critical projects that impact the company in a verifiable way.

B. Innovation Score:

Rating 1 (Low): Responsible for implementing specifications developed by senior engineers and product managers. Does not develop products.

Rating 2 (Middle): Develops incrementally innovative ideas that can be successfully patented. Does not take leadership of developing new products, features and lines of business.

Rating 3 (High): Develops patentable ideas that lead to breakthrough improvements. Comes up with ideas to expand their projects and may also have a reasonable free-hand in developing and executing on them.

C. Leadership Score:

Rating 1 (Low): Tech, design or architectural lead of a small team/project, but could not have direct reports.

Rating 2 (Middle): Be able to mentor engineers in the team, giving technical guidance, code reviews, and ultimately be able to take responsibility of delivering small projects end-to-end on production.

Rating 3 (High): Leads complex initiatives and technically drives teams towards implementing and productionizing them. Promotes professional growth and development inside and outside the team. Actively takes steps to increase technical excellence across the organization.

5.B Additional Information about Job Opening for a Full-Stack Software Engineer

The position of software engineer will involve work on a specific project critical to a start-up's needs with opportunities to change projects and teams as the software engineer grows. Engineers are required to be multifaceted, display successful leadership abilities, and be enthusiastic to tackle new and challenging problems.

Responsibilities may include:

- Design, develop, test, deploy, maintain, and improve software
- Manage individual project priorities, deliverables, and deadlines
- Collaborate with other specialists in development teams
- Analyze and improve efficiency, scalability, and stability of various system resources

Minimum Qualifications:

- BA or BS degree in Computer Science or related technical field
- Experience with one or more general purpose programming languages including but not limited to: Java, C/C++, C#, Objective C, Python, JavaScript, or Go
- Experience working with two or more from the following: web application development, Unix/Linux environments, mobile application development, distributed and parallel systems, machine learning, information retrieval, natural language processing, networking, developing large software systems, and/or security software development
- Working proficiency and communication skills in verbal and written English

N.6 Recruiter Online Evaluation Form

Brief Feedback on 8 Job Applications

We expect this task to take a maximum of two hours. We will ask you six sets of questions about the candidates, so you may want to keep any research open until the end of those questions. We will then ask you seven quick background questions about yourself.

On each page, this form will save your responses as you enter them, you can come back to a page before hitting "Submit" and complete the evaluations of candidates in any order. However, once you have hit "Submit" on any particular page, you will not be able to revise your responses any more.

Click "Submit" to get started.

Submit

Brief Feedback on 8 Job Applications

Please review the candidate information in the packets we sent you. In the table below, tell us:

- Which candidates do you suggest interviewing?
Then, for each candidate, assume he/she was interviewed and passed.
- Suggest a take-it-or-leave-it annual base salary offer for each candidate.
Remember that it's worth paying more to make sure that higher quality candidates say yes. But we also are on a budget and do not want to overpay.

Please provide a salary offer even for candidates you do NOT suggest interviewing. Someone else at our company might suggest interviewing these candidates. If they pass, we will use your input for what salary to offer them.

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. You may enter \$0, but only if that is what you truly intend, otherwise please enter a non-zero value.

	Candidate's	Candidate's	Interview		Your take-it-or-leave-it offer	Notes/Comments (Optional)
	First name	ID #	Yes	No	Amount In Dollars	
1	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
2	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
3	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
4	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
5	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
6	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
7	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>
8	<input type="text"/>	<input type="text"/>	<input type="radio"/>	<input type="radio"/>	<input type="text"/>	<input type="text"/>

We will interview all candidates you suggest. We may also interview additional candidates. Interviewers will not know who suggested each candidate. We will consider your salary responses only after we finish interviewing and deciding which candidates to pursue.

Submit

Brief Feedback on 8 Job Applications

If you could make a single take it or leave it offer, which candidate would you want to make it to?

Notes/Comments (Optional)

As a reminder, here are your answers to the previous questions.

First Name	ID	Interview	Salary
Andrew	1111	Yes	150,000
Christoper	2222	No	150,000
Emily	3333	Yes	150,000
Jessica	4444	No	150,000
Joshua	5555	Yes	150,000
Samantha	6666	No	150,000
Sarah	7777	Yes	150,000
Tyler	8888	No	150,000

Submit

Brief Feedback on 8 Job Applications

Occasionally, candidates make us a take-it-or-leave-it offer (rather than us making one to them). Our policy is to take these proposals seriously. We will either accept or reject the candidate's offer, and the outcome is then final. We do not make counter offers or consider the candidate's follow-up offers. Candidates know this is our policy.

Please note: We do not feel that candidates who make us an offer are necessarily better or worse than those who wait for the offer.

We need your suggestions about how to respond in these settings. Please select the maximum offer we should accept for each candidate in this situation. For your reference, a table of the salary offers you suggested earlier can be found at the bottom of this page.

	Candidate's First name	Candidate's ID #	Maximum take-it- or-leave-it offer we should accept	Notes/Comments (Optional)
1	Andrew	1111	<input type="text" value=""/>	<input type="text"/>
2	Christopher	2222	<input type="text" value=""/>	<input type="text"/>
3	Emily	3333	<input type="text" value=""/>	<input type="text"/>
4	Jessica	4444	<input type="text" value=""/>	<input type="text"/>
5	Joshua	5555	<input type="text" value=""/>	<input type="text"/>
6	Samantha	6666	<input type="text" value=""/>	<input type="text"/>
7	Sarah	7777	<input type="text" value=""/>	<input type="text"/>
8	Tyler	8888	<input type="text" value=""/>	<input type="text"/>

As a reminder, here are your answers to some previous questions.

First Name	ID	Interview	Salary
Andrew	1111	Yes	150,000
Christopher	2222	No	150,000
Emily	3333	Yes	150,000
Jessica	4444	No	150,000
Joshua	5555	Yes	150,000
Samantha	6666	No	150,000
Sarah	7777	Yes	150,000
Tyler	8888	No	150,000

While we will interview all candidates you suggested, we may also interview additional candidates. For your responses, assume all candidates are interviewed and have passed. We will consider these responses only after we finish interviewing and deciding which candidates to pursue.

Submit

Brief Feedback on 8 Job Applications

Suppose we were interested in hiring each of the candidates.

Tell us:

A salary so low, they'd be only 5% likely to take it?

A salary they'd be just as happy taking or rejecting (50% likely to accept or reject).

A salary so high, we think they're very likely to take it? (95% likely to say yes)?

We are looking for 3 distinct salary values for each candidate. For example, for an entry level cashier job potential answers might be: there is a 5% chance they would take \$7.25, there is a 50% chance they would take \$9, and a 95% chance they would take \$15. Use your best judgement.

	Candidate's		Salary			Notes/Comments (Optional)
	First name	ID #	5% likely to yes	50% likely to say yes	95% likely to say yes	
1	Andrew	1111	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
2	Christopher	2222	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
3	Emily	3333	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
4	Jessica	4444	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
5	Joshua	5555	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
6	Samantha	6666	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
7	Sarah	7777	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
8	Tyler	8888	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

While we will interview all candidates you suggested, we may also interview additional candidates. For your responses, assume all candidates are interviewed and have passed. We will consider these responses only after we finish interviewing and deciding which candidates to pursue.

Submit

Brief Feedback on 8 Job Applications

We're interested in your opinion of each candidate's competing offers - please enter your assessments of their potential competing offers below

	Candidate's First name	Candidate's ID #	Number of Compelling Offers (not including current job)	Reason for Competing Offers	Notes/Comments (Optional)
1	Andrew	1111	<input type="text"/>	<input type="text"/>	<input type="text"/>
2	Christopher	2222	<input type="text"/>	<input type="text"/>	<input type="text"/>
3	Emily	3333	<input type="text"/>	<input type="text"/>	<input type="text"/>
4	Jessica	4444	<input type="text"/>	<input type="text"/>	<input type="text"/>
5	Joshua	5555	<input type="text"/>	<input type="text"/>	<input type="text"/>
6	Samantha	6666	<input type="text"/>	<input type="text"/>	<input type="text"/>
7	Sarah	7777	<input type="text"/>	<input type="text"/>	<input type="text"/>
8	Tyler	8888	<input type="text"/>	<input type="text"/>	<input type="text"/>

Submit

Brief Feedback on 8 Job Applications

Given your experience with this type of work, do you think that a job candidate like the ones you reviewed here would mis-report his/her most recent salary (even if we could verify past salary later)?

Yes, he/she may mis-report

No, he/she would not mis-report

Submit

Brief Feedback on 8 Job Applications

You said a job candidate might mis-report his/her most recent salary, if a candidate stated his/her current salary was \$90,000, what do you think this candidates true most current salary is?

Submit

Brief Feedback on 8 Job Applications

Thank you for your work. We just have a few very short additional questions we would like you to answer

Upwork Profile URL:

How long have you been doing this type of work?

Less than 3 months

3 months to 1 year

1 to 3 years

3 to 10 years

Over 10 years

How often do you provide salary input during hiring?

Never

Sometimes

About half the time

Most of the time

Always

How did you make judgments on the salary related questions? Please check all that apply.

Used my previous experience with salaries in this setting

Looked up salaries on a website like payscale.com, glassdoor.com or others

Spoke with others who are familiar with salaries for software engineers

Other

Submit

Brief Feedback on 8 Job Applications

The questions below are optional questions on background for statistical purposes.

I identify my gender as:

Male

Female

Other

I identify my ethnicity as (select as many as apply)

Asian

Black/African

Caucasian

Hispanic/Latin American

Native American

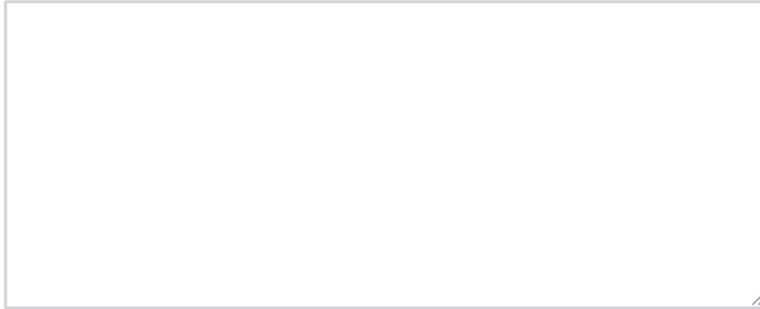
Pacific Islander

Other

Submit

Brief Feedback on 8 Job Applications

Do you have any suggestions to improvements in how we solicit advice from recruiters? This may include technical issues you had, or bigger picture things like other questions we should ask about evaluating candidates?

A large, empty rectangular box with a thin black border, intended for the user to provide feedback. A small cursor icon is visible in the bottom right corner of the box.

Submit

Brief Feedback on 8 Job Applications

We thank you for your time spent taking this survey.
Your response has been recorded.

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.
- Bessen, James E, Chen Meng, and Erich Denk**, “Perpetuating Inequality: What Salary History Bans Reveal About Wages,” *Working Paper*, 2020.
- 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**, “Salary History Bans and Unravelling,” *Working Paper*, 2022.
- Dee, Thomas S**, “A teacher like me: Does race, ethnicity, or gender matter?,” *American Economic Review*, 2005, 95 (2), 158–165.
- 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.
- Hansen, Benjamin and Drew McNichols**, “Information and the Persistence of the Gender Wage Gap,” *National Bureau of Economic Research Working Paper*, 2020.
- 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.
- , **Azeem M Shaikh, and Yang Xu**, “Multiple hypothesis testing in experimental economics,” *Experimental Economics*, 2019, 22 (4), 773–793.
- Mask, Joshua**, “Salary History Bans and Healing Scars from Past Recessions,” 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.

Roussille, Nina, “The central role of the ask gap in gender pay inequality,” *Working Paper*, 2020.

Sinha, Sourav, “Salary History Ban: Gender Pay Gap and Spillover Effects,” *Available at SSRN 3458194*, 2020.

Sran, Gurpal, Felix Vetter, and Matthew Walsh, “Employer Responses to Pay History Inquiry Bans,” *Working Paper*, 2020.