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

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 and participants at the MEER Conference, the Advances in Field Experiments Conference, the American Economic Association Annual Meeting, the Innovation and Institutions 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, the Federal Reserve Board, Florida State University, George Mason University, Nuffield College (Oxford), the 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) for valuable feedback and comments. 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.

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
JEL No. C90,J70,M50

ABSTRACT

We study how salary history disclosures affect employer demand by using a novel, two-sided field experiment featuring hundreds of recruiters reviewing over 2000 job applications. We randomize the presence of salary history questions as well as candidates' disclosures. We find that 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, however they are negative signals of value (net of salary), and thus yield fewer callbacks. Male wage premiums are regarded as a weaker signal of quality than other sources (such as the premiums from working at higher paying firms, or being well-paid compared to peers). Recruiters correctly anticipate that women are less likely to disclose salary history at any level, and punish women less than men for silence. In our simulation of bans, we find no evidence that bans affect the gender ratio of callback choices, but find large reductions in gender inequality in salary offers among candidates called back. However, salary offers are lower overall (especially for men). A theoretical framework shows how these effects may differ by key properties of labor markets.

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

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

Laura K. Gee
Economics Department
Tufts University
8 Upper Campus Road
Braker Hall, Medford, MA 02155
laura.k.gee@gmail.com

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

1 Introduction

Since 2016, 21 states and 21 jurisdictions in the United States have banned employers from asking job applicants about their salary histories.[1] The goal of these laws is to ameliorate historical inequalities relating to gender, race/ethnicity, or other protected characteristics. However, the information encoded in historical salaries is complex. Salary history is provided voluntarily, and employers may anticipate strategic aspects of disclosure. When revealed, salary histories may signal subtle information about candidates' hidden characteristics and outside options. Disclosure behaviors may differ between groups of candidates, requiring complex inferences by employers.

How employers use salary history signals — and what inferences they make in their absence — is the subject of this paper. We study how salary history disclosures affect employer demand and how salary history bans shape hiring and wages. Our conceptual model shows how these effects depend on the characteristics of the labor market. We then measure the key underlying properties – and directly measure the effects of disclosures and bans – in a novel, large-scale field experiment.

Our analysis is focused on three central questions: First, what inferences do employers make about candidates who refuse to disclose salary history? Second, among candidates who do disclose it, what do employers infer from higher or lower salary amounts? Finally, how do employers combine these inferences into decisions about whom to hire and how much to pay them? Our conceptual framework shows how these answers depend on two key properties: i) How correlated salary history amounts are with workers' quality and outside offers, and ii) whether a high previous salary is correlated with low value (net cost) to the employer – or vice versa.

To answer these questions empirically, we developed a novel, two-sided field experiment. Using an intermediary firm, we assumed the role of an employer and hired hundreds of recruiters to make decisions about 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 they were not aware that our job applicants were fictitious. On the demand 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. Among candidates who disclose them, we also vary the levels

[1] See https://www.hrdive.com/news/salary-history-ban-states-list/516662/ for the most up-to-date list.

of salaries, mimicking real-world, male-favoring gender wage gaps.

We then measure recruiters' callback recommendations and salary offers for all candidates. By examining different bargaining protocols, we also infer recruiters' maximum salary offer (e.g., willingness-to-pay, or WTP) for all candidates as well as their beliefs about each worker's competing offers. To contextualize these findings, we present the results of a large survey of the U.S. workforce that measured the prevalence of salary history questions and candidates' responses to them.

Our field experiment yields three main 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. We find that women are punished less for silence. Our survey suggests a possible reason: women at all salary levels are less comfortable disclosing their salary, which makes their choice to disclose less informative about value.

Our second set of results has to do with the amounts attached to the historical salaries disclosed. On average, higher salaries increase recruiter beliefs about candidate quality and competing offers. Every extra $1.00 increase in disclosed salary increases i) employer WTP by $0.65, ii) beliefs about the median competing offer by $0.77, and iii) the salary offer by $0.68. We find that recruiters *discount* extra dollars given to men through the gender wage gap. An extra dollar given to a male candidate (because of the gender wage gap) raises WTP by only $0.42 ($0.48 for the salary offer, and $0.62 for the median competing offer).

By contrast, recruiters regard an extra $1.00 coming from working at a high-wage firm, or being well-paid within a firm's internal distribution as being far more informative (each extra dollar increasing WTP by $0.64 to $0.70, or $0.20 per dollar more than the male bonus). These sizes suggest that recruiters *anticipate* overpaid men, interpreting their inflated salaries as being less indicative of value. However, the effects are also significantly above zero; recruiters are discounting the male premium by about half, but far from 100%.

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

The third set of results from our field experiment detail how recruiters synthesize the inferences above into decisions. Effects on salary offers are straightforward: disclosing workers— particularly those with high salaries— receive higher salary offers. This is especially true for men, whose disclosures are perceived as more informative. By contrast, we find *negative* effects on who is recommended for a callback. Disclosing workers – especially those with high salaries – are less likely to be recommended 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. The reasons for this are theoretically grounded and appear in our data: callback decisions are not a function of worker quality only, they must also incorporate expected costs. This phenomena arises directly from the empirical results described above: salary history disclosure increases beliefs about outside offers more than candidate quality. At some point, outside offers are so competitive that employers' margins are squeezed. Choices of whom to call back mirror those for our measure of expected employer surplus.

For these reasons, we see that non-disclosers (and low-disclosers) are called back more often, but are offered lower average salaries. This pattern has gendered consequences: men are less likely to be chosen when they disclose. On the margin, employers interpret higher male salaries as constituting a prohibitively high price tag. Women's disclosures have less of an effect on whether they are recommended for a callback; this is consistent with their disclosure choices being less informative. We find similar results on the amounts: lower amounts disclosed by women increase their odds of being recommended.

As a whole, our results highlight important trade-offs. In our setting, disclosures— and disclosing higher amounts— increase the level of salary offers but decrease callback rates. Normative implications about the ban thus depend partly on how policymakers prioritize these competing effects and the associated risk preferences. Additional compensation is obviously useful; however, additional outside offers can also give workers leverage and options (even at a lower salary rate). They may also allow workers greater choice in finding good matches on non-salary dimensions.

In our final section, we use the results of our field experiment to evaluate salary history bans and the design choices embedded in these policies. We show that salary history bans reduce gender inequality in offered wages conditional on callback. They achieve this in part by reducing the salaries of all workers and particularly men. The ban has little impact on callback inequality. Some states adopt a "partial" ban that allows employers to ask for salary history after making an initial offer; we show that these design choices are less effective at reducing gender inequality.

This paper provides four main contributions, which we detail in the next section. First, we contribute a novel application of disclosure theory to policies intended to "blind" employers and reshape statistical discrimination, and we present field experimental evidence about how disclosures are understood by employers. 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 tests, we find that recruiters incorporate expectations about gender differences into their decisions. Although their anticipation does not fully eliminate gender disparities, it does reduce inequalities significantly below their original levels.

Third, we extend the literature on "price as a signal of quality" to a labor market setting, and we introduce the idea of "price as a signal of competition." Our theory and empirics show how price revelations influence both perceptions of candidate quality and perceptions of outside offers. In our setting, price is a stronger signal of competition than of candidate quality (particularly for gender-related price differences).

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 salary history disclosures propagate through the hiring process. We gather a rich collection of theoretically-motivated outcomes from the field about both wage setting and candidate selection. To our knowledge, this is the first paper to randomize a government-mandated policy in an audit experiment. Our two-sided design creates multiple avenues for studying discrimination more broadly in future research.

The rest of the paper is organized as follows. The remainder of this section describes three related literatures and our contribution to each. Section 2 provides a brief background about the practice of asking job candidates for salary histories. Section 3 describes a theoretical 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, and we use them to simulate the effects of bans in Section 8. Section 9 concludes with a brief discussion.

1.1 Literature

A large and well-developed literature documents and explores the causes of the income gaps across genders and races.[2] Our paper is also about the persistence of negative wage shocks in the labor market (Kahn, 2010; Oyer, 2008). Within this larger body, our research is related to several strands of prior research.

Voluntary Disclosure and Statistical Discrimination. Our work relates to policies to "blind" decision makers—for example, by banning employers from seeking candidates' credit scores, criminal records, or drug-test results (Clifford and Shoag, 2016; Agan and Starr, 2018; Doleac and Hansen, 2020; Wozniak, 2015) or through gender-blinding résumés (Behaghel et al., 2015; Åslund and Skans, 2012).

In many settings, this form of blinding requires cooperation from the supply side, which can override blinding through voluntary disclosure. This introduces the potential for unraveling (Viscusi, 1978; Grossman and Hart, 1980; Milgrom, 1981; Grossman, 1981; Jin et al., 2015). In many models, lack of disclosure may be viewed as a negative signal of quality, leading to full revelation. Empirically in markets with quality disclosures such unraveling is not always observed (Dranove and Jin, 2010; Mathios, 2000). In surveys, women are less likely to disclose income than men (Goldfarb and Tucker, 2012; Agan et al., 2020; Cowgill et al., 2021), and female software engineers are less likely to disclose their skills than are their male counterparts (Murciano-Goroff, 2017). There is a key difference between salary history and other disclosure games: with salary histories, it's possible for disclosures to be "too high" as well as "too low." By contrast, in most disclosure games, the state variable is monotonically related to the principal's objective ("more is always better"). Our conceptual framework demonstrates this potential theoretically, and our data show this phenomenon in a survey and field experiment.

Gender differences in job search. A variety of prior research documents female job candidates as being less aggressive in job search, having lower propensity to enter competitive environments (Niederle and Vesterlund, 2007; Flory et al., 2015), self-promoting less (Exley and Kessler, 2019), asking for lower salaries from employers (Roussille, 2020), or being less willing to disclose their salary (Goldfarb and Tucker, 2012; Cowgill et al.,

[2] See Goldin (1990); Solnick (2001); Eckel and Grossman (2001); Blau and Kahn (2006); Thilmany (2000); Blackaby et al. (2005); Takahashi and Takahashi (2011); Moss-Racusin et al. (2012); Mazei et al. (2015); Ponthieux and Meurs (2015); Juhn and McCue (2017); Rozada and Yeyati (2018).

2021).[3]

In our applied setting, we propose a microfoundation for this difference: the gendered differences in behavior are the byproduct of gendered differences in the psychological costs of disclosure. These cost differences are correlated with gender, but are uncorrelated with latent characteristics such as talent. Although 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; Niederle, 2015). In our framing, the act of disclosing enters workers' utility function directly.[4]

We then draw out the implications for the demand side. We show that anticipating gender differences is useful for the employer, and that treating the same negotiation signals differently by gender is necessary to update beliefs accurately (Spence, 1973; Fryer Jr, 2007). Our model connects gendered differences in negotiation behavior to theories of voluntary disclosure, employer learning, and unraveling.

We then measure these responses empirically in a field experiment. We present novel evidence on the demand side's response to gender differences in negotiating behavior. Our empirical results suggest that corporate 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 reported by men. Although recruiters' anticipation is insufficient to fully eliminate gender disparities, it does reduce inequalities significantly below their original levels.[5]

Our paper sheds new light on how employers interpret gendered differences in self-promotion and negotiation aggression. We show that these are not only interpreted as a signal of candidate quality, but also as signals about the competitive landscape. In short,

[3]Other examples include the propensity to apply for a job given the number of other applicants (Gee, 2019), the choice to disclose skills (Murciano-Goroff, 2017), the decision to reapply to an employer following a rejection (Brands and Fernandez-Mateo, 2017), and the choice to negotiate wages (Laschever and Babcock, 2003; Biasi and Sarsons, 2020).

[4]Exley et al. (2020) similarly studies negotiation costs by gender, including indirect costs of unsuccessful negotiations.

[5]Murciano-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.

our results suggest that 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 other employers view the candidate.

Price as a signal of quality. Our paper is also related to prior literature about using price as a signal of quality. This idea was originally applied to consumer products in industrial organization, but is less-developed in labor settings. Seminal papers by Wolinsky (1983) and Milgrom and Roberts (1986), study price as a signal of quality,[6] 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 the amount and level of competing offers in a worker's search. These beliefs can affect wages through a separate, non-quality channel. We formalize this notion in our model and relate it to common-versus-private value labor markets. The interaction of these signals is critical to choices about revealing (versus concealing) historical prices.

We then provide direct experimental evidence on how historical prices affect employer beliefs about both quality and outside competition for workers. In most of our results, we find a greater role for price as a signal of competition (versus as a signal of quality). In particular, our recruiters interpret gender-related price differences as signals of competition, rather than of quality.

Audit Methodology. Methodologically our work is related to recent innovations in correspondence audit methodology (Bartos et al., 2016; Kessler et al., 2018; Kline and Walters, 2019; Avivi et al., 2021; Cowgill and Perkowski, 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." In most traditional audit studies, the researcher cannot control the characteristics of the employers. Our two-sided design allows manipulation of employer characteristics in a randomized way. To our knowledge, this is the first paper to randomize a government-mandated policy in an audit experiment. This

[6] Roussille (2020) adapts the Wolinsky (1983) model to a labor market.

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

two-sided design creates multiple avenues for studying discrimination more broadly in future research.

Salary History Bans. A nascent literature directly studies salary history bans introduced in the past decade. 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, 2019; 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., 2019).

Our field experiment uses recruiters for corporate jobs, and is focused on the microeconomics underlying these policies (voluntary disclosure, unraveling, and from 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 salary history bans across jurisdictions.

Given these research goals, we use an experimental paradigm that delivers access to detailed variables in the hiring function (including WTP, beliefs about competing offers and expected employer surplus). These variables are typically missing from administrative data sets and do not directly appear in the other experiments, but are useful for understanding the mechanisms and trade-offs around the ban.

2 Background: Salary History Questions

Survey evidence suggests that up to 43% of job applicants are asked about salary history during job search (Hall and Krueger, 2012; PayScale, 2017; Barach and Horton, 2021; Agan et al., 2020; Cowgill et al., 2021). In our own survey, the most common method of inquiring about salary history was on job application forms (in writing). Among workers who were asked, 45% were asked this way (Agan et al., 2020; Cowgill et al., 2021). Written salary history questions on job application forms are so common, in fact, that some jurisdictions explicitly address the practice in the text of their bans.[8] Job interviews (34%)

[8] These jurisdictions include the states of Virginia, for public jobs, and New York, as well as municipal-level bans in Atlanta, Ga.; Jackson, Miss.; Suffolk County N.Y.; and Richland County, S.C. See `https:`

were the second most common context of the question.

In August 2016, Massachusetts adopted the first ban on salary history questions (effective July 2018), thus becoming the first jurisdiction to adopt such a policy. As of August 2021, 21 states and 21 local jurisdictions have adopted some form of salary history bans.[9] In 2019, a federal salary history ban passed the House of Representatives (and passed again in 2021).[10] President Biden has agreed to sign if the federal ban passes the Senate, and has issued an executive order to ban salary history questions in federal agencies' hiring practices.[11] These laws vary in their details, but nearly always prohibit oral or written questions about salary history, even if the questions are posed as voluntary or optional.[12]

However, applicants under the bans are still permitted to *voluntarily* and *without prompting* disclose salary history information. In most jurisdictions, employers are allowed to use or confirm voluntarily disclosed information.[13] In addition, the laws do not ban employer questions about a candidate's hopes, wishes, or expectations around compensation.

The popularity of these bans masks enormous heterogeneity in the *designs of salary history bans*. For example, some bans apply only to either state or jurisdiction agencies (public employers), and others apply to all employers in the jurisdiction (both public and private). Our experiment addresses heterogeneity in the design of the ban. Some jurisdictions ban employer from asking **until an initial offer has been made** (but *can ask afterwards during negotiation of the initial offer.*).[14] This design is focused on giving all candidates a common "floor" to begin negotiating. By contrast, other jurisdictions ban asking

//www.hrdive.com/news/salary-history-ban-states-list/516662/.

[9]See https://www.hrdive.com/news/salary-history-ban-states-list/516662/ for the most up-to-date list. This list includes the two states that have banned local jurisdictions from salary history bans.

[10]2019: https://thehill.com/homenews/house/436121-house-passes-paycheck-fairness-act, 2021: https://www.cnbc.com/2021/04/16/what-the-paycheck-fairness-act-could-mean-for-women-and-the-pay-gap.html

[11]https://www.law360.com/employment-authority/articles/1401105/biden-targets-pay-equity-with-salary-history-executive-order

[12]The goal of the bans is expressed in the title of the laws: "The Act to Establish Pay Equity" in Massachusetts, or "An Act Concerning Pay Equity", in Connecticut. On the day the Illinois law was announced, one of the sponsors said, "I am proud to stand with our new governor today as he takes strong, immediate action to close the gender pay gap and move towards pay equity[.]" https://www2.illinois.gov/Pages/news-item.aspx?ReleaseID=19609

[13]Some 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.

[14]This 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.hrdive.com/news/salary-history-ban-states-list/516662/ and https://www.ebglaw.com/news/new-york-state-releases-guidance-on-salary-history-ban/.

at any point in hiring, both when making callback decisions and when setting wages. This includes California, and New York (from 2020 to the present, following a revision). This approach aims to eliminate salary history not only from the floor, but from the negotiation altogether. This latter distinction about when employers can (if ever) ask about salary history is an important component to understanding the potential impact of salary history bans. It will come up again in our simulations in Section 8.

3 Conceptual Framework

We now offer a simple theoretical framework of hiring and wage setting in which salary history disclosures (or lack thereof) are a key input. This framework will offer insight into when and how salary history disclosures will appear, and how they shape hiring decisions, thus providing guidance for our empirics. Our framework starts from a standard unraveling model (Viscusi, 1978; Grossman and Hart, 1980; Milgrom, 1981; Grossman, 1981), with a few key adaptations to our setting.

Unraveling is a classic result that suggests that if information is verifiable and the costs of disclosure are small, market forces will compel all participants to disclose information voluntarily about their own characteristics. The intuition behind unraveling is simple: An audience of buyers will view all non-disclosing parties as equals, and thus high-quality agents have an opportunity to distinguish themselves through volunteering information (i.e., sharing good news about themselves). When they do, buyers' lower their expectations about the quality of the remaining non-disclosers. This logic proceeds iteratively until all parties disclose, or until all non-disclosed salaries can be deduced.

To adapt this model to our setting, we make three adjustments to the baseline model. We summarize these and discuss their implications below, and we use the remainder of this section to specify our implementation of these changes.

1. First, the audience for the disclosure (employers) is making two, interrelated decisions for each candidate: whom to call back, and how much salary to offer. Consistent with prior research (Postel-Vinay and Robin, 2002; Dey and Flinn, 2005; Cahuc et al., 2006), we conceptualize these two decisions as a function of the employer's beliefs about the candidate's value to the employer (WTP), and about the candidate's outside options.

2. Second, the audience (employers) does not have preferences directly over the vari-

able that candidates can disclose or not (in this case, salary history). Instead, the disclosed variable is a noisy signal used to update beliefs about the variables above (WTP and outside offer). Because it is a noisy signal, disclosures may not be fully informative. Below, we discuss some reasons why salary history may be more (or less) informative in labor markets.

3. Finally, all candidates have other observable characteristics (such as gender). Observable characteristics are partially correlated with the hidden variable (salary history). As a result, the employer can make an informed guess about the candidate's hidden value even without a disclosure. This makes disclosure particularly impactful for candidates whose hidden values are rare (given observables).

   In addition, candidates have hidden *costs of disclosing*. The presence of disclosure costs means that silent workers contain a mixture of candidates with unattractive salaries (that were strategically withheld), and candidates with high costs. Costs can be correlated with an observable characteristic (in our case, gender). Our survey suggests that women have higher disclosure costs. The implication is that recruiters may know that a silent woman is more likely to have high disclosure costs, and not necessarily a low value, and that a silent man is more likely to have a low value.

Adaptations No. 2 and No. 3 are straightforward. In the remainder of this section, we specify the form of the two decisions (callback and salary level) referenced in No. 1 which the employer makes after receiving the message and updating beliefs. We show that the effects of disclosing depend on three features of the market.

First, how correlated are salary histories with workers' latent characteristics (such as their quality and competing options)? This feature determines whether disclosures are noisy or informative. Second, does the employer surplus for hiring a worker[15] tend to be positively (vs negatively) correlated with their previous salary? This affects the direction of employers' updates in response to salary information. Finally: How much are workers willing to trade-off the probability of getting a callback at all, with higher/lower salary offers (conditional on a callback)? This answer also affects how employers interpret disclosures. In the section below and in our empirics, we show how these issues shape disclosure.

[15] By this, we mean the worker's value minus the compensation necessary to retain that worker.

3.1 Preliminaries: Players, Utilities and Notation

The two players in our model are an employer and a job applicant. The employer's utility comes from money and the productivity from hiring a candidate. The employer chooses to call a candidate back (or not) with a take-it-or-leave-it (TIOLI) salary offer. If hired, a job applicant yields a stream of utility payments to the employer over multiple time periods ($v_1 + v_2 + ...$) and salary payments to the worker ($s_1 + s_2 + ...$) from the employer (these payments may be zero after a worker quits or is fired). The net present values of these streams are v and s. Making a callback costs the employer c.

The applicants' utility comes from salary payments and the disutility of work. Each applicant has a set of k outside offers. The outside offers are the net present value of the salaries offered by other employers. The k offers are drawn from a distribution H, and includes the candidate's current position if currently employed.[16] A candidate's best outside option, η, is the first order statistic from the k draws from H. k and H are known to the employer, but the realizations are not. For any given salary offer s by the employer, the candidate accepts with probability $\Pr(s > \eta)$.

Given a joint distribution of beliefs $F(v, \eta)$ about the candidate's value and outside offers, the employer can calculate a TIOLI offer $s^\star$:

$$\begin{aligned} s^\star &= \operatorname*{argmax}_s \mathbb{E}\left[\underbrace{\mathbb{1}(s > \eta)}_{\substack{\text{Whether candidate}\\ \text{accepts}}} \cdot \underbrace{(v - s)}_{\substack{\text{Net value of}\\ \text{employment,}\\ \text{if accepts}}} \right] \\ &= \operatorname*{argmax}_s \iint \underbrace{\mathbb{1}(s > \eta)}_{\substack{\text{Whether candidate}\\ \text{accepts}}} \cdot \underbrace{(v - s)}_{\substack{\text{Net value of}\\ \text{employment,}\\ \text{if accepts}}} \cdot \underbrace{f(v, \eta)}_{\substack{\text{Joint}\\ \text{probability}}} \, dv d\eta \end{aligned} \tag{1}$$

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

[16] Additional draws come from a combination of the active search for other employers or from employers' active recruitment of the candidate.

$$\mathbb{E}\Bigg[\underbrace{\mathbb{1}(s^\star > \eta)}_{\substack{\text{Whether candidate}\\ \text{accepts the optimal}\\ \text{TIOLI offer } s\star}} \cdot \underbrace{(v - s^\star)}_{\substack{\text{Net value of}\\ \text{employment,}\\ \text{if accepts the}\\ \text{optimal TIOLI}\\ \text{offer } s^\star}}\Bigg] = \iint \underbrace{\mathbb{1}(s^\star > \eta)}_{\substack{\text{Whether candidate}\\ \text{accepts the optimal}\\ \text{TIOLI offer } s\star}} \cdot \underbrace{(v - s^\star)}_{\substack{\text{Net value of}\\ \text{employment,}\\ \text{if accepts the}\\ \text{optimal TIOLI}\\ \text{offer } s^\star}} \cdot \underbrace{f(v,\eta)}_{\substack{\text{Joint}\\ \text{probability}}}\, dv d\eta > \underbrace{c}_{\substack{\text{Fixed cost}\\ \text{of a callback}}} \tag{2}$$

Note that higher beliefs about v increase the employer's payoffs from giving a callback and for offering a generous TIOLI amount. However, higher beliefs about outside options increase $s^\star$, but decrease returns of sending a callback (unless beliefs about v also increase).[17] The employer cannot justify a callback to an expensive employee, unless the employee's contributions make it worthwhile. Neither v nor η can be directly observed. However, the employer can estimate these from observable characteristics, the choice to disclose salary history, and the amount if disclosed. We now turn to how these beliefs are formulated as a function of disclosures.

3.2 Updating Beliefs Using Salary History

Candidates have a hidden salary variable h that can either be disclosed or not. h can be verified if disclosed, and thus the candidate cannot lie.[18] The candidate's action space is $\{h, \emptyset\}$. Working backwards, we now turn to the candidate's choice to disclose. Candidates receive utility from offers and from paying (or avoiding) a cost to disclose.

If the candidate receives no callback, the payoff from the callback is η (his privately-known outside option). If he gets a callback, his payoff is α for receiving a callback at all, and an additional β for every dollar above his outside offer (η). If the amount is below his outside offer, he receives only α. This could be positive because an offer —even at a low salary— can potentially be useful to a job seeker.[19] α and β are known to the employer. The candidate's utility is thus:

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

[18]This is the standard requirement in the disclosure literature.

[19]We have three specific scenarios in mind. First, a low initial salary may still be useful as a bargaining chip in renegotiations with one's current employer. Second, a worker may gain utility from non-pecuniary aspects of the job. Stern (2004) finds that scientists choose a 30% pay cut to be able to do science. 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 low offer may be useful.

$$u = \underbrace{\eta}_{\text{Outside offer}} + \underbrace{\alpha \cdot Callback}_{\text{Additional payoff from getting a callback}} + \underbrace{\beta \cdot \mathbb{1}(s > \eta)(s - \eta)}_{\text{Additional payoff from getting a callback with a salary above the outside offer}} \tag{3}$$

Since the employer is making salary choices based on Equation 1 and callback decisions according to Equation 2 (and the candidate knows this), we can rewrite the candidate's utility in terms of the employer's offer $s^\star$, which is in turn a function of the employer's beliefs about the candidate's value and outside offer (Equation 1):

$$u = \underbrace{\eta}_{\text{Outside offer}} + \alpha \underbrace{\mathbb{1}\left(\mathbb{E}[\mathbb{1}(s^\star > \eta)(v - s^\star)] > c\right)}_{\text{Does the employer believe the candidate meets the requirements for a callback? (Equation 2)}} + \beta \underbrace{\mathbb{1}\left(s^\star > \dot{\eta}\right)}_{\text{Does candidate get an offer above outside option } \dot{\eta}} \underbrace{\left(s^\star - \dot{\eta}\right)}_{\text{Amount above outside offer}} \tag{4}$$

... where $\dot{\eta}$ represents the employer's true, privately known value of η.

To understand the candidate's disclosure strategy, we must specify how h is correlated with the candidate's value v and his or her outside offers η. The joint distribution of η, v and h is publicly known. While this distribution could in theory take any shape, two aspects are important:

Assumption 1 (Informativeness)**.** v (value) and η (outside option) are correlated with h, and $E[v|h]$ and $E[\eta|h]$ are both weakly increasing in h.

This assumption means that h (the hidden salary variable), is a noisy signal of value (v) and outside option (η). These correlations could be high in settings such as a common value labor market, or zero in settings where firms value the same workers very differently (i.e., when there are match-specific sources of productivity, Jovanovic 1979a,b; Lazear 2009).[20]

Assumption 2 (Monotonicity)**.** The expected value of the worker above his or her outside option ($E[v - \eta|h]$) — is weakly monotonic in h. This simply means that on average, the slope of h on v is steeper than h and η (or vice versa).

In principle, the monotonicity of $E[v - \eta|h]$ mentioned in 2 could either be increasing or decreasing. High h candidates could, in theory, deliver the highest surplus.[21] In this

[20]In principle, both correlations could be negative for a given employer.
[21]One setting where this appears to be true is professional sports, where leagues with salary caps prevent

case, as h increases, the worker would be paid an increasingly small fraction of the value they provide. This would be true if a high salary offer were more highly correlated with valuations than outside offers. From the perspective of disclosure strategy, candidates disclosing high amounts receive both higher callback chances and higher salary offers (conditional on a callback) than low disclosers.

However, the relationship could also go in the opposite direction: salary histories could be negatively correlated with employer surplus. High h candidates could deliver lower surplus. As h increases, workers are paid a larger fraction of their value. This would happen if high salaries increase outside options faster than the benefits. At some point, outside options may catch up with value, leaving the candidate too expensive to warrant a callback. In this scenario, high h candidates face a trade-off: disclosing may decrease their callback chances, but increase the salary amount conditional on callbacks. Is the trade-off worthwhile? It depends on the value of receiving a callback at a low amount (α) compared to the value of additional dollars (β).

Which of the two scenarios describes reality is an empirical question we address with our experiment. As we discuss later, our results more closely resemble the second scenario in which high h candidates are increasingly bad deals: a higher salary history increases beliefs about η more than beliefs about v, and thus at some point candidates become too expensive to callback. Some candidates thus face a trade-off between high salaries and a callback that depends on α versus β (the utility weight placed on callbacks versus additional dollars).

3.3 Choices to Disclose and Unraveling

The candidate can now make a disclosure decision. Disclosure is costly and requires a cost m that is privately known to the candidate, and drawn from a distribution F_m that is publicly known. The candidate can anticipate the employer's posterior beliefs about v and η following a disclosure. With these posteriors, the candidate can anticipate whether she will receive a callback from the employer and at what price using Equations 1 and 2. The candidate can then use Equation 3 to determine the payoff from a disclosure.

If he does not disclose, then the standard unraveling logic proceeds: the employer assumes that the candidate had an unattractive salary and proceeds as if the salary history

employers (teams) from paying the full value of superstars. As such, superstars with the maximum permissible contracts under the cap are regarded as cheap compared to the value they bring, and thus a good use of dollars. Salary minimums (Horton, 2017) may also affect employer preferences between candidates.

was bad. This makes disclosure more attractive, even for those with unattractive salaries. As described above, this proceeds until all candidates disclose, or until the benefit of disclosing is equal to the cost m.

Where our story differs from the standard unraveling story is around what constitutes an "unattractive" salary. Because high h candidates are bad deals, salaries could be "too high" if candidates place a high weight on getting a salary offer at all (high α relative to β). As β becomes larger compared to α, the candidate cares only about high salary unconditional on a callback. The strategy is to disclose if high, hide if low. The standard unraveling logic proceeds. Silence is assumed to be a mixture of low salaries and high costs of disclosure (high m realizations from F_m). How low a salary should one disclose? Until benefits stop exceeding m.

If α becomes larger than β, the candidate cares only about getting an offer. The best strategy is to disclose if low, hide if high. The standard unraveling logic proceeds *in reverse*: low salaries should be disclosed, and high salaries should be hidden. Silence is assumed to be a mixture of *high* salaries and high costs of disclosure. How high a salary should one disclose? Again, until benefits stop exceeding m.

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.

3.4 Observable Characteristics and Costs

As described at the beginning of this section, candidates differ on observable characteristics. Our framework above easily accommodates this; the joint distribution of h, v, and η can differ by observable characteristics. As a result, employers' beliefs and reactions to candidates who behave the same way could differ, even if for candidates with an identical h.

We also allow costs m to vary by observable characteristics. Costs of disclosing play a key role in our theory by regulating how far unraveling proceeds, and thus how informative silence is. For candidates with observable characteristics i with high average disclosure costs, unraveling will be more limited. Rather than signaling the worst salary, silent types will contain a mixture of relatively good and bad types (depending on how high costs are).

At first glance, disclosing salary history would appear to be cheap using a pay stub, bank statement or offer letter. However, workers report feeling *psychological costs* of disclosing. In our survey a majority of workers expressed that they were uncomfortable answering salary history questions by a potential employer.[22] In addition, these costs are not evenly distributed. Across multiple studies, women are less comfortable disclosing their salaries than men (Goldfarb and Tucker, 2012; Agan et al., 2020; Cowgill et al., 2021). In our own survey, women were about 6 percentage points more likely to agree with statements about feeling discomfort with disclosing than men. This is salient for our conceptual framework and experiment, because it suggests that a woman's silence will contain less information about their underlying characteristics.

3.5 Implications for the Ban

The ban affects only the employer's ability to ask questions. Most unraveling models do not feature a question, as candidates can voluntarily disclose. We interpret employer questions (or other prompts) as lowering m_i (the costs of disclosing); and thus bans raise this cost for all individuals.

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 1. In 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.

[22]Surveyed workers were asked to rate how much they agreed/disagreed with the statement "I am fundamentally uncomfortable answering the salary history question when asked by a potential employer',' on a seven-point scale from strongly disagree to strongly agree. Answers of five or above were coded as agreeing in some form. More than half, 51%, agreed in some form with 21% saying they strongly agreed. Some 17% of respondents strongly agreed with the statement, "I can't think of anything that would make me fully comfortable providing my salary history (when asked by a potential employer)," and 44% expressed some form of agreement.

4 Empirical Setting

To apply our conceptual framework to data, we examined disclosure in the labor market for software engineers. We focus on this industry for reasons outlined in Appendix A: broadly speaking, the characteristics of this industry help place our results in the context of wider labor market trends. Our experimental design utilizes the managerial practice of delegating recruitment decisions to recruitment specialists, whom we hire to make decisions about candidates in our two-sided audit study. In this section, we provide an overview of this setting and review institutional details that motivate our experimental design.

4.1 Outsourced Recruiting

In the past two decades, the delegation of recruiting has become widespread (Cowgill and Perkowski, 2021). Although comprehensive data on firm hiring practices is not available (Oyer et al., 2011), surveys estimate that between 40% and 50% of U.S. firms outsource some portion of their recruiting process or plan to do so soon, and that as many as one-third of jobs in the United States are filled through outsourced recruiting.[23] Firms either hire individual recruiters on a temporary, contract basis, or they outsource recruiting work entirely to a third party organization. A large industry known as the recruitment process outsourcing (RPO) industry exists to support this practice. The recruitment process outsourcing industry has been growing steadily since 2000, and is growing faster than other sectors.[24] The subjects in our field experiment are recruiters employed in this industry.

[23] https://www2.staffingindustry.com/content/download/272268/9900694/StaffingTrends_Free_190301.pdf https://www.workforce.com/2019/01/24/recruitment-process-outsourcing-providers-think/.

[24] According to Cowgill and Perkowski (2021), outsourced recruiting has been growing at 4 to 10 times the rate of overall U.S. employment since 2000, using the BLS' Occupational Employment Statistics (https://www.bls.gov/oes/home.htm) Additionally, the number of recruiters in employment services has more than doubled, while U.S. employment overall grew only 15% from 2002 to 2009 according to the U.S. Economic Census. The growth in outsourcing recruiting practices is driven by many factors including increases in the popularity of Internet job search, numbers of applicants, demand for employee screening (Autor, 2001; Cappelli et al., 1997), and returns to selective hiring (Acemoglu, 1999, 2002; Levy and Murnane, 1996).

4.2 Recruiter Work

Industry benchmarking reports show that recruiters perform a wide variety of tasks for employers.[25] Surveys indicate that nearly all firms that use outsourced recruitment ask their recruiters to screen applications, and over 95% of recruiters have been asked to provide input on the salary of potential hires.[26] Even before the COVID-19 pandemic, over 80% of outsourced recruiting was performed remotely.[27] Recruiters in this industry are typically told to avoid searching on the Internet for information on job candidates, as this can violate employment law.[28] Our study design closely mimics each of these attributes.

4.2.1 Structured Evaluations:

Recent surveys of recruiters by Jobvite,[29] Monster.com,[30] and Black et al. (2020) indicate that assessing candidates using structured and/or quantitative evaluation is typical and expected to grow. An industry organization, The Talent Board, found that 71% of employers used pre-employment assessments and selection tests in 2019.[31] Westin (1988) reviews the history of quantitative assessment and formal job testing. Employment testing grew out of twentieth-century personality psychology, which has a long tradition of quantitative assessment scales.

4.3 Recruiter Incentives

The use of monetary incentives is widespread in recruiting. The monetary incentives a recruiter faces can take many forms and are often quite complex.[32] Employers often use a formula for bonuses that rewards recruiters for selecting candidates who are eventually

[25] See https://www.shrm.org/ResourcesAndTools/business-solutions/Documents/Talent-Acquisition-Report-All-Industries-All-FTEs.pdf for example.

[26] See https://staging.kornferry.com/media/sidebar_downloads/Measuring-Up-A-new-research-report-about-RPO-metrics.pdf and Analysts (2017) and our own survey responses from subjects in this experiment.

[27] *Staffing Industry Analysts*, RPO Market Developments, December 2017.

[28] For 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.

[29] https://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

[31] https://www.thetalentboard.org/report/the-2019-talent-board-north-american-candidate-experience-benchmark-research-report-now-available/

[32] Recruiter.com states that the "complexity of recruiter commission plans tends to rival both ontological arguments and mortgage refinancing documents." See https://www.recruiter.com/i/recruiter-commission-plans/

hired and rated as high performers (either on the job, or during the selection process).[33]

If recruiters were rewarded for selecting candidates who are hired, they may be tempted to forward candidates indiscriminately. Recruiter incentives therefore can include measures to induce selectivity, so that the employer does not waste time on candidates unlikely to match. In some cases, firms impose an explicit cap on the number of candidates a recruiter can forward. In settings where an employer is seeking all qualified candidates, firms may use an explicit monetary penalty for forwarding candidates who are a bad fit.[34] Beyond monetary incentives, recruiters face reputational incentives. Even without explicit monetary rewards or penalties, recruiters know an employer wants them to forward a curated list of candidates, and give useful advice about fit and compensation. Explicit monetary schemes reinforce these implicit reputational concerns.

5 Experimental Design

To examine the impact of salary disclosure on candidate outcomes, we implemented a two-sided audit study. We hired a real recruiting workforce to screen (fictitious) candidates on behalf of a (fictitious) firm. This design allows us to vary both characteristics of the candidates *and* characteristics of the firm. To our knowledge, this design has been utilized in only one other paper (Cowgill and Perkowski, 2021), and we extend the design to fit our research agenda. On the candidate side, variation comes from experimental changes to candidates' gender, salary, and whether that salary is disclosed at all during the application process. On the employer side, there is variation in whether the firm asks for salary information, and the number of candidates who disclose their salary history. Each recruiter's job was to review eight candidates and provide feedback about each to the firm, with questions and format detailed below.

[33] According to industry reports 60% of performance pay is measured by the number of candidates hired. See https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/rewarding-recruiters-for-performance.aspx. A particularly well-documented example of this practice is the U.S. military's recruiter bonus programs (Condren, 1997; Asch, 1990). The Bureau of Labor Statistics' National Compensation Survey (NCS) found that 43% of human resource specialists received performance-pay in the first quarter of 2020 (Makridis and Gittleman, 2020). This percentage likely understates the extent of incentives for recruiters; the NCS definitions of performance pay include incentive payments that "follow an explicit formula" and discretionary spot bonuses (Gittleman and Pierce, 2015). Other sources of performance-based rewards that are excluded include permanent wage increases, promotions, and the possibility of new assignments or repeated business (for contracted recruiters).

[34] HR Magazine https://www.shrm.org/hr-today/news/hr-magazine/pages/1103hirschman.aspx

5.1 Our Recruiting Workforce

We hired 256 recruiters to evaluate 2,048 job applications for our field experiment. These recruiters are our experimental subjects. We identified recruiters who are typical of those hired by companies through the recruitment outsourcing industry and engaged them in a natural way for this industry. Using a large and popular outsourcing platform, we identified and contacted professional recruiters as discussed in Appendix C. 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, with an additional bonus to be described shortly.

The human resource workers in our field experiment were similar to those in the U.S. as a whole.[35]Each subject was assigned to one of the experimental conditions (described below) using the randomization procedure in Appendix B.1.

5.2 Our Recruiting Task

Recruiters who accepted our offer were given three documents to complete their task: 1) detailed instructions describing the job they were hiring for, 2) eight one-page PDF job candidate applications, and 3) an online structured evaluation form to provide feedback about the candidates. Examples of these are provided in Appendix H.

As described in Section 4.2, there is wide variation in the tasks recruiters perform. Our goal was to design a task and incentive scheme that allowed us to answer our research questions while fitting naturally into our setting. In addition to their hourly rate, we offered recruiters an additional bonus described in Appendix C. In the main text of our communications with recruiters, we described in simple, non-technical language how the bonus worked. This was likely sufficient for many of our recruiters. However, the details of the bonus including an explicit formula were available in the FAQ portion of the same PDF document available in Appendix H.

Recruiters were required to answer each question about all candidates before proceed-

[35]According to the BLS in 2018 Human Resource workers across all industries were 69.7% female, 10.5% black, and the median hourly wage was $29.01 across all industries, and $41.93 in the software industry.[36] 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 G2). The BLS does not report demographic characteristics of industry×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.

ing to the next question for any candidate. Recruiters were not told the specific items on the structured evaluation beforehand. Each group of questions was on a separate page which did not allow recruiters to revise previous answers.[37]

Below, we detail the items in the structured candidate evaluation. In this manuscript, we describe recruiter choices in this evaluation using theoretical terminology. However, these concepts also appear in typical hiring settings in less formal language, without models or academic jargon. Recruiters and hiring managers often consider questions like, "How much will this worker contribute?'," "How much are we willing to pay?'," or "What will their competing offers look like?" Our experiment measures these common economics concepts without the use of economics terminology. A full copy of the structured evaluation can be seen in Appendix H.6.

5.2.1 Primary/Incentivized Assessments The three variables below (callback, offer and WTP) were the first items evaluated by recruiters. These items were also explicitly used in each recruiter's bonus formula (and many counterpart bonus formulae in the real world).[38] The recruiters therefore faced monetary incentives to set these numbers correctly.

5.2.1.1 ***Callback*** Like a traditional audit study, we observe whether each recruiter recommends a candidate for a callback. Recruiters were aware that they should feel free to suggest multiple callbacks (as many as they deemed a good fit) rather than limit themselves to filling a single position. This is common in high-tech labor settings featuring high demand for qualified workers. We conceptualize this as the callback decision modeled in Equation 2 of our framework.

5.2.1.2 ***Salary Offers*** Recruiters made a take-it-or-leave-it (TIOLI) salary offer for each job candidate. We conceptualize these as the $s^{\star}$ choice modeled in Equation 1 of our framework. 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 observe a TIOLI offer even if the recruiter did not believe the candidate should be called back.[39]

[37] If we later found an error in an item (e.g., a recruiter typed in letters in an item that required numbers), we then contacted the recruiter and sent them a new link to revise that item.

[38] See Appendix H.5 for details of our specific bonus formula and how it was communicated to recruiters.

[39] Recruiters 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

The offer and callback decisions were on the first page of our evaluation. After these responses were submitted, recruiters could not revise them.

5.2.1.3 ***Willingness to Pay (WTP)*** 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.[40] We conceptualize these as the v (expected value) that appears in our framework. By reporting a threshold, we essentially observe the recruiter's value for the candidate directly. Any value below their true value is a dominated strategy (Becker et al., 1964).

5.2.2 **Additional Assessments** After the questions above, recruiters made additional assessments that were not explicitly mentioned in the bonus criteria. Although recruiters did not face formal contractual incentives to formulate these numbers precisely, they still face reputational incentives for accurate answers.

5.2.2.1 ***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 to mean 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.

5.2.2.2 ***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." In our conceptual framework, this roughly corresponds to k, the number of outside offers drawn from the candidate's outside offer distribution. Prior research suggests that outside offers increase the bargaining power of the candidate (Blackaby et al., 2005), and that employed workers

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."

[40] Recruiters 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.

rarely receive more than one job offer at a time when searching.[41] Recruiters also state whether competing job offers from each candidate would come from the candidate's own search efforts, or from rival employers' search efforts. These variables shed light on why the levels of competing offers (η) are as high or low as they are.

5.2.3 Composite Outcomes We take some of the main items we observe from recruiters and combine them into the following outcomes of interest.

5.2.3.1 ***Surplus*** We label the value of the candidate as measured by willingness to pay (WTP) minus the suggested salary offer as the surplus a recruiter assigns to a candidate. Our theoretical framework Section 3 describes how this can be interpreted as the margin the employer would get from this candidate. This corresponds to $\mathbb{E}[v - s^\star]$ in our conceptual framework; and it appears in Equations 2 and 4. It is also the topic of the monotonicity requirement in Section 3.2 (Assumption 2), which states that surplus is an increasing or decreasing function of salary history. As discussed in 3.2, 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).

5.2.3.2 ***Outside Offer Range*** We define the outside offer range as the difference between the 95th percentile and the 5th percentile of the outside offer distribution. This gives us information about how wide ranging the recruiter's beliefs are about a candidate's outside offers. Larger ranges indicate more diffuse beliefs. This is akin to the variance of the belief distribution of η (the best outside offer).

5.2.3.3 ***Probability of Accepting TIOLI Offer*** This variable represents how likely the recruiter believes the candidate is to accept the take-it-or-leave-it salary proposed by the recruiter at the start of the form (described in 5.2.1.2). In our conceptual framework, it is represented as $\mathbb{E}[s^\star > \eta]$ and appears in Equations 1, 2 and 4. To estimate this probability, we fit a logistic curve through the reported 5th, 50th and 95th percentiles of the outside offer distribution.[42] We used this fitted model to predict where the salary offer made by the recruiter falls on that curve. In our theoretical framework from Section 3 we discuss

[41] Faberman et al. (2017) find that only 29.1% of employed workers who are looking for work receive at least one offer per month.

[42] Similar exercises using probit curves produced results with correlations of 0.99 with the logistic approach.

how the probability of accepting an offer relates to whether the salary offer is above or below a candidate's outside offer.

5.2.3.4 ***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 $\mathbb{E}[(s^\star > \eta)(v - s^\star)]$ and appears in Equations 1, 2 and 4.

5.3 Experimental Manipulation

Having described our outcomes, we now turn to experimental manipulations. Our two-sided audit study entails experimental manipulation on both the employer side *and* the candidate side.

5.3.1 Employer Side On the employer side, we manipulated whether the job applications reviewed by a recruiter asked the candidates for their previous salary, or not. In our survey of workers, we found that 44.8% of workers who were asked about their previous salary were asked on the job application (Cowgill et al., 2021). Asking about salary history on the job application form is so common that several states specifically ban salary history questions on the job application form (in addition to verbal questions). Our question asked for the applicant's annual base salary at their current or most recent job.

5.3.2 Candidate Side In addition, 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.

5.3.2.1 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.[43] Our hypotheses about gender are also related to the disclosed salary amounts (described next).

[43] The names were: Andrew, Christopher, Joshua, Tyler, Emily, Jessica, Samantha, and Sarah.

The second randomized biographical detail is about low- and high- wage firms. To randomize this, we utilized a list of the 13 biggest employers of software engineers from Monster.com and Indeed.[44] This included substantial variation in median salaries for software engineers who recently graduated.[45]

In order 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 avoid suspiciously identical candidates. Appendix D lists these details.

Using this procedure, we created 32 candidate biographies 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" to increase the efficiency of the experiment.[46]

5.3.2.2 Salary History Disclosures Candidates' salary history disclosures (or lack thereof) were also randomized. In packets where candidates were asked salary histories, candidates who disclosed that information completed the question on the form by entering a number on the line. In packets where candidates were *not* asked, candidates who disclosed it used an optional field for "Additional Skills and Information."[47]

The amount of disclosed salary is also randomized (among those who disclosed). We designed these disclosures to be consistent with the candidate's biography, but to include some random variation (conditional on the biography). For each candidate's current em-

[44] 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/.

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

[46] A biography consists of the specific combination of name, gender, previous and current employers, job titles and descriptions, dates of employment, undergraduate education, and skills.

[47] This section also included information about programming skills the applicant had. For disclosure without a prompt, the candidate would add a sentence such as "Current Base Salary: $X yearly." We randomized whether the statement appeared before or after the additional skills, as well as the language of this statement slightly.

ployer, we looked up the distribution of salaries for software engineers at the candidate's location and job level using Payscale.com.[48] 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.

Importantly, we also built in a gender wage gap. These gaps were based on real world gaps in this industry. Female candidates were given a salary that was 85% of the male salary at the same company/location.[49] Despite the gender wage gap built among our candidates, most variation in the disclosed wages can be explained by previous employers and 75th/25th fixed effects. As such, we present targeted, disaggregated specifications in our results that isolate the effects of gender-related wage differences. Appendix D contains additional details about how we assigned salaries to candidates.

5.4 Subtreatments

In theory, all of the attributes above could be randomized independently at the candidate level. However, 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. Either 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 disclosed and sent a separate packet to another recruiter flipping the candidates' disclosures. This form of clustering allowed us to measure how candidates' disclosure choices affected each other through spillovers.

Finally, we clustered candidates' disclosure amounts. As previously mentioned, candidates' disclosure amounts were randomized (for those who disclosed). We clustered

[48] Data 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).

[49] To approximate realistic gender gaps in salaries, we analyzed data from the 2015 American Community Survey (ACS) on individuals in computing jobs and found women earned around 85% of 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.

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 others candidates' outcomes).

For the full details of our clustered randomization scheme, see Appendix B. The clustered randomization produced 22 subtreatments, where a subtreatment is a combination of {asked, not} × {all disclose , half disclose, other half disclose, none disclose } × {all high amounts, half high +half low, other half high + half low, all low amounts}.[50] In some specifications below, we control for subtreatment fixed effects.

5.5 Randomization Procedure and Balance

Because we randomize both sides of the market, we check for randomization balance in both candidate and recruiter characteristics. Appendix D.1 and B.1 show our candidate and employer/prompting manipulations were uncorrelated with each other (partly by construction) or with the characteristics of the assigned recruiter. Appendix B contains the details of how recruiters were assigned to treatments.

6 Specifications

Our goal is to estimate how candidate salary disclosures— both their existence and amounts— affect hiring outcomes. Our specifications come mostly in one of three forms, outlined by the equations below.

6.1 Specification 1: Effects of Disclosure vs. Silence

$$y_{i,j} = \beta_1 SalaryDisclosed_{i,j} + \beta_2 SalaryHistoryAsked_j + \nu_i + \beta_3 [RecruiterControls_j] + \epsilon_j \tag{5}$$

where i indexes candidates and j indexes recruiters. Outcomes $y_{i,j}$ are the assessments given to candidate i from recruiter j. Whether the recruiter saw applications with a salary history question is controlled for by $SalaryHistoryAsked_j$. Our bundle of recruiter controls ($RecruiterControls_{i,j}$), includes the gender, race, experience level, and hourly rate of

[50]Our total number of treatments is less than $2 \times 4 \times 4 = 32$ because in cases where no candidates disclose, amounts are irrelevant.

the assigned recruiter (which were balanced by design, section B.1). ν_i signifies candidate biography fixed effects (see section 5.3.2).

We also extend this specification in two main ways: First, we introduce candidate gender interactions with the $SalaryDisclosed_{i,j}$ terms. The main effect of candidates' gender is absorbed by the biography fixed effects (ν_i). Second, we replace $SalaryDisclosed_{i,j}$ with a vector $\{DisclosedLowSalary_{i,j}, DisclosedHighSalary_{i,j}\}$ where $DisclosedLowSalary_{i,j}$ is a dummy variable for whether the applicant disclosed a salary at the 25th percentile of his or her current firm salary distribution and $DisclosedLowSalary_{i,j}$ is whether that applicant disclosed at the 75th percentile of his or her current firm salary distribution.

6.1.1 Interpretation

We interpret the $SalaryDisclosed_{i,j}$ coefficient (β_1) in Equation 5 as the average effect of disclosing. This represents whether disclosers are (on average) paid more or less than silent types. A positive coefficient indicates that recruiters believe that above average types have selected into disclosing. When we include instead whether the candidate disclosed low (25th percentile of within firm salary distribution) or high (75th percentile of within firm salary distribution), this helps identify where in the distribution silent candidates are presumed to lie.

6.2 Specification 2: Effects of Disclosure Amounts

The next specification measures how the amount the candidate discloses impacts outcomes by adding $AmountDisclosed_{i,j}$ terms to Equation 5.

$$\begin{aligned} y_{i,j} =& \beta_1 SalaryDisclosed_{i,j} + \beta_2 SalaryDisclosed_{i,j} \times AmountDisclosed_{i,j} \\ &+ \beta_3 SalaryHistoryAsked_j + \nu_i + \beta_4[SpilloverControls_{i,j}] \\ &+ \beta_5[RecruiterControls_j] + \epsilon_j \end{aligned} \tag{6}$$

We set $AmountDisclosed = 0$ for candidates that did not disclose (their overall impacts are captured by "$SalaryDisclosed = 0$"). As described in Section 5.4, our experiment was designed to study potential spillovers between candidates' disclosures. Our regressions thus include a set of $SpilloverControls_{i,j}$. We treat these as control variables and do not report spillover coefficients in our tables in this paper.[51] Spillover controls are detailed in Appendix E.

[51] Our main interest in this paper is the average direct effects of one's disclosures on one's own outcomes.

6.2.1 Interpretation The slope coefficient on $AmountDisclosed_{i,j}$ (β_2) represents how much outcomes $y_{i,j}$ are influenced by the amount disclosed. Coefficients with larger magnitude (steeper slopes) indicate greater influence on a recruiter's evaluation. Greater magnitudes are consistent with recruiters relying more on the disclosure amount—i.e., updating or changing their beliefs about the candidate based on the information in the disclosure amount. In the opposite extreme (in which the $AmountDisclosed_{i,j}$ slope is flat), recruiters do not incorporate new information from the disclosure amount, and recruiter behavior does *not* change at all as the amount varies. We will use this specification to measure the direction of our monotonicity assumption (Assumption 2) in Section 3.2. This requires that the employer's surplus be increasing (or decreasing) in salary history amounts, or that the salary history has a greater slope with valuation than outside offers (or vice versa).

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

Finally, our results also include a third specification that focuses on the subcomponents of $AmountDisclosed_{i,j}$. Here, we unpack the reasons why candidates in our experiment are paid differently, and we show how recruiters' reactions depend on these differences. As discussed in our candidate variation section (5.3.2), candidates are paid differently for three reasons: 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, and thus report differently even conditional on other factors.

A plausible null hypothesis is that recruiters will treat each source of pay variation equally. Each candidate presents a single number (the sum of these factors), not individual components. Nothing in our experiment alerts recruiters to *any* possible reason for pay variation, much less the three above. Nonetheless, recruiters could possibly anticipate these reasons. Certain employers are well-known for paying well. Recruiters who know a candidate works for Apple may adjust expectations upwards about prior salary. Recruiters who know that women are typically paid less may adjust expectations about salary history downward. As discussed earlier, the only source of variation that cannot be anticipated is the within-firm variation, which we designed to be uncorrelated with any observable feature.

In order to measure these effects in our data, we utilize the following identity:

$$AmountDisclosed_{i,j} = OverallAverage + FirmOffset_i + GenderOffset_i + WithinFirmOffset_{i,j} \tag{7}$$

Each candidate's disclosure amount is the sum of an overall average across all candidates ($OverallAverage$, a constant), plus a firm-specific offset for the candidate's employer ($FirmOffset_i$: some firms pay higher or lower to everyone on average), a gender offset ($GenderOffset_i$, penalizing women and favoring men), and a within-firm offset ($WithinFirmOffset_{i,j}$, representing pay variation for the same job). The sum of these is (by definition) equal to the total amount disclosed. These relationships flow directly from our procedure for creating salary disclosure amounts (Section 5.3.2).

We use this definition to replace $AmountDisclosed_{i,j}$ in Equation 6 with the subcomponents in Equation 7. This leads to our third specification:

$$\begin{aligned} y_{i,j} =& \beta_1 SalaryDisclosed_{i,j} + \beta_2 SalaryDisclosed_{i,j} \times FirmOffSet_i \\ &+ \beta_3 SalaryDisclosed_{i,j} \times GenderOffset_i \\ &+ \beta_4 SalaryDisclosed_{i,j} \times WithinFirmOffset_{i,j} \\ &+ \beta_5 SalaryHistoryAsked_j + \nu_i + \beta_6[SpilloverControls_{i,j}] + \beta_7[RecruiterControls_j] + \epsilon_j \end{aligned} \tag{8}$$

This decomposed regression allows us to obtain separate coefficients for each source of variation in salary, and compare them. We use these coefficients to study how recruiters evaluate an extra dollar of historical salary differently, depending on *why* the salary is higher or lower. Earning +\$1 extra because a worker is at a high-wage firm may evoke a different response than if the same +\$1 came from being a star engineer at a lower-wage firm. The extra \$1 may evoke a different response if it was awarded for being male.

Because prior employer, gender and within-firm salary are randomly assigned, these separate coefficients can each be interpreted causally. We call these the heterogeneous effects of disclosing an extra \$1. We can separate these effects from the direct effects of being male and employed at a high-wage firm, because the extra +\$1 effects are only present in our observations where candidates disclose.

As with the slopes in Equation 6, the magnitude of these coefficients represents how informative they are. If recruiters anticipate any differences by using biographical features (Apple and/or male workers are typically paid more), this anticipation would *re-*

duce the informational content of the disclosure, and push these coefficients toward zero. Of course, gender and employer differences may *not* be fully anticipated, and thus we could also find nonzero results.

7 Results

By design, our experiment collected information about not only final outcomes (such as whom to make an offer to and what the salary amount offered should be), but also on the inputs into those decisions such as the maximum salary offer the firm should accept (willingness-to-pay) and information about employer perceptions of the candidate's outside offers. We start by showing how these measures of inputs and outcomes are related in ways our conceptual framework portrayed. We then examine how a candidate's silence or disclosure impacts employer perceptions about these inputs, and how these change with the amount disclosed. Finally, we explore the implications for callback and salary decisions.

In Table G3 we show correlations between outcomes and recruiters' WTP for candidates and their beliefs about competing offers— relationships in our data modeled by our theoretical framework. As expected, recruiters' TIOLI offers to candidates are positively correlated both with willingness-to-pay for the candidate, and with recruiters' beliefs about the candidates' competing offers. The TIOLI/WTP relationship is stronger than the relationship between TOILI and outside offers. However, as our theoretical framework showed, callback decisions can be positively correlated with willingness-to-pay, but negatively correlated with competing offers. Higher outside offers increase the price that firms must pay to attract a candidate and thus reduce the firm's surplus for selecting the worker. Although these results are descriptive, they are in line with our portrayal in the conceptual model.

7.1 Effects of Disclosure vs. Silence

7.1.1 Overall / Average Effects Our recruiters believe non-disclosers have lower quality (as measured by WTP) and lower outside options, and are given lower salary offers, compared with the average disclosing worker. This pattern can be seen in our basic descriptive statistics in Table 1: recruiters' WTP, assumptions about outside offers, and offers are lower for candidates who do not disclose in column (3) compared to those who do in column (2).

In the odd columns of Table 2, we model these outcomes using Equation 5. When a candidate discloses, recruiters increase willingness-to-pay by about \$6,800 (6.46% over the mean of non-disclosers). Likewise, disclosure increases perceptions of outside offers by \$8,400 an 8.8% increase over the mean for non-disclosers. These results suggest a level of sophistication by recruiters, who appear to anticipate positive selection into disclosing. Disclosure choices have a larger effect on recruiters' beliefs about candidates' competing offers, and a smaller effect on recruiters' willingness-to-pay – but both are affected in the same direction. These increases in WTP and outside offers lead to an increase in offers of about \$7,300 or a 7.5% increase in the amount offered to those who disclose their previous salary. Our results in Table 2 suggest that recruiters make inferences about and assign outcomes to silent workers that are worse than the average candidate within the candidate's $\{$current employer $\times$ job title $\times$ gender$\}$ wage distribution. In Column 7 we see that disclosure, perhaps unsurprisingly, reduces uncertainty about the level of outside offers— recruiters' perceptions about the spread between an offer a candidate would accept with a 5% probability and one accepted with a 95% probability are compressed by nearly \$9,000 (32% over mean spread of \$28,500).

The results above focus on the average effect of silence versus disclosing, but our design allows us to decompose the effect into low versus high earners within each firm. In Table 3, we show results from our extension of Equation 5, splitting Salary Disclosure into disclosing a low (within current firm) salary versus a high (within current firm) salary. We find that recruiters infer that silent candidates' hidden salaries are at or just slightly below the 25th percentile of potential outcomes (given the candidates' observables). Workers below this percentile are actually better off silent. However, despite the negative inferences we document, we also see that recruiters do not fully punish non-disclosure as much as would be possible. In fact, some workers who disclose low amounts are regarded as even worse than non-disclosers (they might have been better off staying silent).

"Below the 25th percentile" is a significant discount. However, the theoretical literature on disclosure rationalizes much more punitive discounting: silent workers should be assumed to have the worst possible wage given observables. Why aren't recruiters more punitive? As in other settings (Jin et al., 2015), recruiters may be inattentive or naive about the strategic aspect of non-disclosure.

We find evidence of two additional mechanisms. The first is that salary histories can be "too high" as well as "too low." As discussed previously, this is a key differentiator between salary history and other disclosure games in the literature. How can salary histories be "too high"? Our conceptual framework in Section 3 demonstrates this idea theoretically: high salary histories may signal that a worker's salary expectations may exceed that worker's value. Our survey results find supporting evidence for this: workers are less willing to disclose extremely high salaries.[52] In later sections (7.4), we show this intuition to be correct: our recruiters denied callbacks to workers disclosing the highest salaries. Given this, silent candidates might be interpreted differently. Silent workers may contain a mixture of those 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.

We also find a second new mechanism behind the lack of full punishment: recruiters may believe that disclosure/silence choices are not driven entirely by a candidate's current salary level. Some workers may be inherently uncomfortable disclosing. They may feel that revealing salary history and/or allowing it to be used in hiring is repugnant. Our own survey— and others like it— suggests that women in particular do not like disclosing. Women are about 12 percentage points less likely to disclose than men when not prompted, and women disclose about 1 percentage point less often than men when prompted (Table G12).[53] Our survey results suggest that women are less willing to disclose their salaries, even after controlling for salary, education and other characteristics. Recruiters may anticipate this, and realize that female non-disclosure may be less informative. If women dislike disclosing – whatever their incomes are – then we should find that female non-disclosers would contain a greater variety of types (and not just low-salary types). To recruiters who understand this, the implication is that female non-disclosers may not signal a lower historical salary.

This is exactly what we find. In the even columns of Table 2, we interact disclosing with gender. Our results show that women job candidates are punished *less* for non-disclosure. Recruiters penalize a silent man's WTP by $5,700 more, his outside option $9,200 more, and his offer by $6,900 more more than for a silent woman. Thus, recruiters change their beliefs less about a non-disclosing female, compared to a non-disclosing male. The flip side of this behavior, however, is that the benefit of disclosing is also smaller for women.

From the perspective of our first mechanism, this result is particularly striking. Female candidates— who are paid less— should have a *lower* risk of being "too high" for

[52] See Agan et al. (2020) and Cowgill et al. (2021).

[53] The difference when prompted is not statistically significant.

an employer's budget. Silent women should therefore contain a greater proportion of "too low" candidates, and thus be penalized *more* for silence in beliefs about their value, outside options and offers. However, we find the opposite. This is suggestive evidence that our recruiters are somewhat sophisticated consumers of disclosure information and are aware that women perceive a higher psychic cost of disclosing. Some 75% of our recruiters are themselves female (and 69.7% of recruiters in the economy broadly),[54] and 71% have three-plus years of prior recruiting experience (Table G2). For these reasons, recruiters may understand that non-disclosing women may not be hiding bad information, but may simply be less willing to share their prior salary.

One potential alternative explanation for our results is that recruiters simply misjudged the average market wages for this job. We explore this in the appendix (Section F) and find that lack of knowledge does not seem to be driving our results. This is both because recruiters reacted to lack of information by doing more research about average market wages, and because our results are similar when we create a subset on recruiters who see half their job candidates disclosing.

7.2 Amounts Disclosed

Our second set of results is about the amounts disclosed among those who disclosed. In the odd columns of Table 4, we use Equation 6 to measure how recruiters reacted to an additional $1 of salary a candidate disclosed.

On average, higher salaries increase recruiter WTP for candidates. This is consistent with employers believing that prior salary carries information about worker quality. However, WTP does not increase 1:1 with current salary—in column 1 we see that on average, for every additional $1 of current salary disclosed, WTP increases by $0.65. Column 3 shows that recruiters' expectations about the candidates' outside offers similarly increase with higher amounts. The effects on competing offers is actually more responsive to disclosure than WTP. For every $1 of current salary disclosed, this median increases by $0.77 (compared to $0.65 for WTP, statistically significantly different with $p = 0.002$). In aggregate, salary offers have a stronger relationship with willingness to pay than with outside offers (Table G3). However, *disclosures about salary history* have a bigger effect on competing offers. This is consistent with the idea that salary history contains information about the worker's outside option, even if it is less informative about the worker's quality.

[54] See https://www.bls.gov/ooh/business-and-financial/human-resources-specialists.htm.

This idea—that salary history contains more information about worker quality than outside offers—appeared earlier in our conceptual framework. In Section 3, we showed that if salary history was more informative about outside offers than about value, then high salary history workers would yield the lowest surplus and well-paid workers would face a trade-off between getting an offer and getting a high offer. Our results here foreshadow this result and we will direct evidence in a later section.

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 x Disclosed x Amt Disclosed" are all statistically insignificant and small. Our results here are highly aggregated, and we decompose salary variation into multiple sources using Equation 6 next.

7.3 Heterogeneous Effects of an Extra $1

Variation in current salaries can arise from several sources. Some workers in our experiment work at firms with high or low overall wage distributions. Some workers are on the high or low ends of those internal distributions. Last, some are female, and earn less as a result of the gender wage gap. We now examine whether these variations in salary amounts have identical effects. Table 5 presents the results from Equation 6, decomposing our $AmountDisclosed_{i,j}$ variable into multiple sources of salary variation.

We broadly find that not all dollars are created equally. As predicted in our theory, 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, recruiters could anticipate that a worker at Apple (a high-wage firm in our sample) is paid around a certain average, and/or that men are typically paid more. These are represented as "+$10K from Firm" and "+$10K from Male," respectively, in Table 5. Although it is not always different from the other coefficients in a statistical sense, the slope on the firm's component is nonzero and is the second steepest across most of our regressions, suggesting that recruiters do not fully anticipate average firm wage differences.

The gender slope is the flattest and least informative: for each additional $1 candidates disclose because of the gender wage gap which favors men, recruiters give 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 about quality or the candidate's outside offers. Alternatively, they may feel that men's higher reported salaries are uninformative because they are exaggerated lies.

In fact, the gender slope is particularly small (flat, uninformative) for our measure of perceived worker quality (WTP). For every $1 given for within-firm reasons, recruiters increase WTP by $0.70. However, recruiters only increase WTP by $0.42 for every $1 allocated to men through the gender wage gap (WTP vs. outside option $p = 0.00$). Despite lower average amounts, the gender slopes— including those on candidate quality (WTP)— are still far from zero. Even if recruiters update less from the gender component, they still impact the callback and compensation decisions in our experiment.

The slope for an additional $1 from being male is steeper on our measure of perceived outside options, although still less than the other slopes (and still statistically different from zero). While recruiters are reluctant to believe men's higher salaries are a signal of quality, they are more willing to believe they are correlated with the candidate's outside offer.[55] For all three components of salary variation, the slope on the recruiter's WTP is lower than the slope on beliefs about the candidate's outside options.

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

Until now, our empirical results have been about employer beliefs, and whether recruiters have more favorable (or unfavorable) beliefs about candidates depending on their characteristics and choices. Our final set of results is about how recruiters synthesize inferences into decisions about whom to move forward with hiring and how to compensate them. These decisions are modeled as a function of beliefs in Equations 1 (salary amounts) and 2 (callbacks). As we covered above, the effects on salary offers are straightforward: disclosing workers— particularly those with high salaries— receive higher salary offers.

[55]Higher outside offers could come partly from a candidate's search effort, or from rival companies' search efforts targeting the candidate. When we try to measure the direction of search in Table G8, we find mostly small and insignificant slopes. However, the direction suggests that recruiters believe men have higher outside options through higher candidate-driven search. By contrast, the other slopes are also correlated with greater outside offers. However, recruiters believe these offers originate from rival companies' search efforts, rather than from the candidates'.

Our results about callbacks are more nuanced. On average, our results about callbacks generally go in the *opposite* direction as those on salary offers. Disclosing workers— especially men and workers with high salaries— are less likely to be recommended *at all*. Although they enjoy higher salary offers when selected, they are less likely to move forward in the hiring process.

Our estimates in this section are sometimes underpowered because of the binary nature of the outcome variable. In addition, the relationship between callbacks and amounts disclosed is noisy near the lower end of the salary distribution. Where possible, we use the theoretically-motivated outcomes we collected — employer surplus, probability of acceptance and others — to help shed additional light the story.

7.4.1 Callbacks and Amounts Our strongest evidence about disclosures and callbacks comes from our specifications that include amounts. Across several specifications, disclosing a higher salary reduces callbacks. In Figure 2, we visualize the relationship between amounts disclosed and callbacks using our full data set. As the figure reveals, we find a noisy relationship in the lower ≈15% of the data (below $70K). However in the upper ≈85% of our data, there is a visibly negative relationship.[56]

Tables 9 and 10 present these results as regressions using Equation 6 for callbacks and outcomes theoretically related to callbacks (employer surplus, predicted probability that candidate accepts offer, and probability of acceptance × surplus). Table 9 includes our full sample. Even in our full sample, which includes the noisy results on the low salaries, the relationship is downward sloping (albeit small and statistically insignificant). When we focus on the upper 85% of the data in Table 10, we find the stronger negative relationships visible in Figure 2. 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 9 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

[56] To find the 70K cutoff we ran regressions predicting callback as a function of salary disclosed and the square of salary disclosed from versions of Equation 5 systematically creating subsets for the data from those above 60K, 70K, 80K and so on. The squared term is insignificant after we subset to 70K and above. We also ran a regression predicting callback as a function of salary disclosed from versions of Equation 5 systematically subsetting the data from those above 60K, 70K, 80K etc. We found the coefficient on the amount disclosed is negative and statistically significant for amounts $70K and above. Results are available from the authors upon request. This threshold may have arisen at $70K because of the following statement in our instructions: "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."

by \$0.03. Higher disclosure amounts also reduce employer surplus, as well as beliefs about the candidate accepting the offer. Table 10 also presents gender interactions for our slopes. We find that the amount/callback relationships are generally stronger for men, both statistically and economically. This is in contrast to the results in our earlier section in which female amounts featured a higher slope.

As we showed above, we can disaggregate the effects of salary differences. Table 8 examines callback-related outcomes using Equation 6 (using the full sample). This allows us to measure the effects of a \$1 on callback outcomes, depending on whether it came from the gender wage gap, between-firm variation or within-firm variation. For nearly all of our callback related outcomes and sources of variation in Table 8, we find negative relationships with extra dollars. Randomly assigned higher amounts from any source we study are correlated with lower callback outcomes, although many of these effects are statistically insignificant.

Of these, our results on the gender wage gap (\$10K from being male) are especially precisely 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 destroys \$7,000 in employer surplus. The probability of accepting the TIOLI offer (and its product) also declines. These results show 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 again suggests some level of sophistication by recruiters.

This is not a contradiction of our previous results. To the contrary, it is the natural extension of the findings established earlier and in our conceptual model. As our survey and theoretical results show: callback decisions are not entirely a function of worker quality; they must also incorporate expected costs. As discussed earlier in our paper, salary disclosures may also affect whether a candidate is hired at all (and by whom), versus leaving the position unfilled or hiring a different candidate.

Throughout several tests above, we find that salary history disclosure increases beliefs about outside offers more than candidate quality (Tables 2, 3 and 5). Where could such a process lead? As outside offers outpace candidate quality, eventually employers' margins will be squeezed. As salary expectations rise too high, selecting an expensive candidate becomes a bad economic deal. Employers must pay a premium to keep the candidate out of a rival's hands, but the candidate's (perceived) quality may not warrant it. Our findings about callbacks are logical extensions of our earlier findings.

7.4.2 Silence and the Choice to Disclose We now turn to how silence affects callbacks. In typical models of disclosure, the negative interpretation of silence is theoretically related to higher disclosure amounts signaling better news. We see this pattern in our results in Sections 7.2 and 7.3: Higher disclosure amounts contain good information, and lower disclosures (and thus silence) are interpreted negatively. For callbacks, higher disclosure values appear to signal *worse* news to employers. This raises the possibility that silence would actually have a *positive* effect for callback decisions. We examine this in Table 6 using our full sample, and Table 7 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, discussed above), they are also cheaper to keep away from competing firms and opportunities (Table 2 column 3, also discussed above).

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.

8 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 disclosure, 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 in our theory section (3.5), 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 above. 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: the "full ban" (no salary questions, ever) and the "partial ban" (which allows salary history questions after the initial offer).

8.1 Estimation Strategy

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

8.1.1 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 E.1 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 (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.

8.1.2 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).[57] 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

[57] This 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.

useful to workers if they grant flexibility and/or negotiating power with a current employer.[58]

8.1.3 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. 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.

8.1.4 Regression specifications

To estimate the effect of asking salary history questions with a prompt, we use Equation 9 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.

$$y_{i,j} = \alpha + \beta_1 SalaryHistoryAsked_j + \beta_2 Female_{i,j} + \beta_3 Female_i \times SalaryHistoryAsked_j + \epsilon_j \tag{9}$$

Notice that Equation 9 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

[58] One 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.

hiring process (like a TIOLI amount for those called back).

8.1.5 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 9 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.[59] 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 60, we review a weighting scheme as if everyone were an "always-discloser."[60]

We aim to employ weights for Equation 9 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 9 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).

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

[60]Suppose 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).

8.2 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.[61] The surveys asked respondents about their demographics, and whether they would disclose their salary when asked (or volunteer when not asked).[62] 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.[63] 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 G12 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.[64] 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 occupations.

Using this data, we construct eight cells ({male,female} $\times$ {High Wage Firm, Low} $\times$ {High

[61] We used the company Prolific Academic, `https://www.prolific.co/`. In a short technical report (Agan et al., 2020) and a longer paper (Cowgill et al., 2021), we document this survey instrument and analyze the data included in this paper in more detail.

[62] The 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.

[63] When 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., 2021).

[64] They work in slightly higher paying jobs and industries.

Salary within Firm, Low}). The distribution of disclosure types of all eight groups can be browsed in Table 11. 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 11 with our field experimental data, giving weights to the experimental observations that the survey indicates are more likely. The final columns of Table 11 summarize how much each group gains/loses from disclosing (on average).

8.3 Ban Simulation Results

8.3.1 Ban Effects on Callbacks Table 12 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 in employment or job changing from bans (Bessen et al., 2020; Sinha, 2019; Hansen and McNichols, 2020; Mask, 2020).[65] However, these results seem to contrast with our earlier results in which silence *increased* callbacks (Section 7.4). Table 11 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.

[65] Some 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).

8.3.2 Effects of a Ban on Salary Offers Conditional on Callback

Table 13 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 13, 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 ($p = 0.13$).[66]

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 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 G13). 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.

9 Conclusion

This paper assesses salary disclosure microeconomically. We develop a conceptual framework that adapts the voluntary disclosure literature to hiring and wage-setting in labor

[66] The 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.

markets. We then deploy a novel field experimental design— a *two-sided audit*— that utilizes recruiters to peer inside the black box of hiring. Our design permits us to vary characteristics of job applicants and of firms. Our experiment traces disclosures from messages, to updated beliefs, and from employer inferences to choices that affect candidates' pocketbooks and daily lives. The tools we develop could be used to examine a wide variety of other interventions that alter firm hiring practices.

We specifically look at the effects of salary disclosure on the hiring process. Salary disclosures are of interest because salary history bans have become a popular policy instrument to close gender (and other) wage gaps. Our results would be difficult to obtain from observational data sets where one does not usually observe whether a job candidate has disclosed his or her salary, and where the decision to disclose is endogenous and likely correlated with possible confounders like current salary and gender. Using a field experiment allows us an unprecedented level of observation and control over correlates of disclosure.

We have three main empirical results. First, silent job candidates are believed to be less valuable, have lower outside offers, and so are extended lower take-it-or-leave it offers. Second, disclosing is especially beneficial to those with higher salaries. In settings with a gender wage gap men benefit more from disclosing than women in part through their higher salaries. Additionally, women feel more discomfort from disclosing, and knowing this employers may update less strongly in response to silence from a woman than a man. Both the gender wage gap and other non-wage factors mean the benefits of disclosing accrue mostly to men. Third, unlike other settings where more is always better, at some point a salary disclosure can be so high that it makes a worker unattractive to the firm. This leads to our third result which is that for most people disclosing a higher amount lowers the likelihood of obtaining a callback. Our results highlight the trade-offs of disclosing and disclosing high numbers, which results in higher offers but a lower likelihood of making it to the next step in the hiring process.

Next we pair our field experiment with a survey of U.S. workers to simulate the effects of salary history ban policies on gender gaps. Our survey allows us to estimate how workers will comply, or not, with salary history ban attempts to lower the number of salary disclosures. In our setting with a preexisting gender wage gap, we found that salary history bans have little effect on callback rates for men or women. In our first set of results we found that staying silent was generally harmful for outcomes beyond callbacks. Bans are meant to increase silence about salary history, and so it is no surprise that bans which prohibit ever asking for salary history decrease salary offers for men and women. However, these bans also decrease the gender gap because the harm of silence is greater to men than to women. In our setting, we find that salary history bans obtain their stated goal of increasing equity, albeit harming female workers but harming male workers more.

Our paper leaves key questions for future research: First, we do not study whether employers are "correct" in the inferences we document. Are low-disclosing workers truly worse performers? Are they silent about their salaries for strategic reasons? Do they truly have poorer outside offers? The subjects in our field experiment are experienced recruiters who have incentives to make accurate choices. However, we cannot independently assess their judgments. Any such assessment would be challenging: it would require worker-level performance data, including for rejected candidates (to assess whether they had been rejected mistakenly). For these reasons, the beliefs we document may persist, even if they are incorrect.

In addition, it would require knowledge of employers' subjective map between worker-level behaviors and firm utility. These measurement challenges prevent *actual* employers from assessing their own workers; prior research suggests that employer learning can be slow in practice (Lange, 2007; Kahn and Lange, 2014). According to a 2019 survey by the Society of Human Resources Management, 77% of employers do not measure the quality of hires on a regular basis.[67] Employers may lack corrective feedback, may be inattentive or otherwise fail to learn for behavioral reasons (Mullainathan et al., 2008; Jin et al., 2015).

Second, we do not answer whether workers are holistically better off, or more equal, as a result of the interventions we study. Our results mostly speak to employer demand: whether workers are paid more, given higher salary offers, or given an offer at all. Combining these into a single measure of worker welfare involves trade-offs. Several of our results go in opposite directions, for theoretically-grounded reasons. Revealing a low salary can make a good worker appear cheap, and make a callback more likely. Having any form of outside offer gives workers options and negotiating power. However, appearing cheap may also lower the compensation amounts attached to selections. Different workers and policy makers may prefer to resolve these trade-offs differently. Workers and firms could also compensate for the issues we document on other margins, for example, by changing their search intensity.

Many of the topics addressed in this paper – unraveling, price as a signal of quality (and competition), employer anticipation, and trade-offs – are missing from popular policy dis-

[67] https://www.shrm.org/ResourcesAndTools/business-solutions/Documents/Talent-Acquisition-Report-All-Industries-All-FTEs.pdf

cussion. Our paper aims to build on prior research on these areas and better understand these topics.

References

Acemoglu, Daron, "Changes in unemployment and wage inequality: An alternative theory and some evidence," *American economic review*, 1999, *89* (5), 1259–1278.

__, "Technical change, inequality, and the labor market," *Journal of economic literature*, 2002, *40* (1), 7–72.

Agan, Amanda and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *The Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

__, Bo Cowgill, and Laura Gee, "Do Workers Comply with Salary History Bans? Voluntary Disclosure, Adverse Selection, and Unraveling," in "AEA Papers and Proceedings," Vol. 108 2020, pp. 33–37.

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.

Asch, Beth J, "Do incentives matter? The case of navy recruiters," *ILR Review*, 1990, *43* (3), 89–S.

Ashenfelter, Orley C, Henry Farber, and Michael R Ransom, "Labor market monopsony," *Journal of Labor Economics*, 2010, *28* (2), 203–210.

Åslund, Olof and Oskar Nordströum Skans, "Do anonymous job application procedures level the playing field?," *ILR Review*, 2012, *65* (1), 82–107.

Autor, David H, "Why do temporary help firms provide free general skills training?," *The Quarterly Journal of Economics*, 2001, *116* (4), 1409–1448.

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.

Bartos, Vojtech, Michael Bauer, Julie Chytilova, and Filip Matejka, "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition," *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.

Behaghel, Luc, Bruno Crépon, and Thomas Le Barbanchon, "Unintended effects of anonymous resumes," *American Economic Journal: Applied Economics*, 2015, *7* (3), 1–27.

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," Technical Report, National Bureau of Economic Research 2020.

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: Empirical evidence from the UK academic labour market," *The Economic Journal*, 2005, *115* (501), F81–F107.

Blau, Francine D and Lawrence M Kahn, "The US gender pay gap in the 1990s: Slowing convergence," *ILR Review*, 2006, *60* (1), 45–66.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Beliefs about gender," *American Economic Review*, 2019, *109* (3), 739–73.

Brands, Raina A and Isabel Fernandez-Mateo, "Leaning out: How negative recruitment experiences shape women's decisions to compete for executive roles," *Administrative Science Quarterly*, 2017, *62* (3), 405–442.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, "Wage bargaining with on-the-job search: Theory and evidence," *Econometrica*, 2006, *74* (2), 323–364.

Cappelli, Peter, Steffanie L Wilk et al., *Understanding selection processes: organization determinants and performance outcomes*, Center for Economic Studies, US Department of Commerce, Bureau of the Census, 1997.

Clifford, Robert and Daniel Shoag, "'No More Credit Score': Employer Credit Check Bans and Signal Substitution," 2016.

Condren, Conal, *Encouraging Recruiter Achievement: A Recent History of Military Recruiter*, Rand Corporation, 1997.

Cowgill, Bo, Amanda Agan, and Laura Gee, "Salary History Bans and Unravelling: Psychological Costs and Coarse Beliefs," *Working Paper*, 2021.

— **and Patryk Perkowski**, "Delegation in Hiring: Evidence from a Two-Sided Audit," *Working Paper*, 2021.

Croson, 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.

Dey, Matthew S and Christopher J Flinn, "An equilibrium model of health insurance provision and wage determination," *Econometrica*, 2005, *73* (2), 571–627.

Doleac, Jennifer L. and Benjamin Hansen, "The Unintended Consequences of "Ban the Box": Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden," *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.

Eckel, Catherine C and Philip J Grossman, "Chivalry and solidarity in ultimatum games," *Economic Inquiry*, 2001, *39* (2), 171–188.

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, National Bureau of Economic Research 2017.

Flory, Jeffrey A, Andreas Leibbrandt, and John A List, "Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions," *The Review of Economic Studies*, 2015, *82* (1), 122–155.

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.

Gittleman, Maury and Brooks Pierce, "Pay for performance and compensation inequality: Evidence from the ECEC," *ILR Review*, 2015, *68* (1), 28–52.

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.

Goldin, Claudia, "The gender gap: An economic history of American women," *Cambridge UniversityPress, New York, Estados Unidos*, 1990.

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," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 56–67.

Hansen, Benjamin and Drew McNichols, "Information and the Persistence of the Gender Wage Gap: Early Evidence from California's Salary History Ban," *National Bureau of Economic Research Working Paper*, 2020.

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," 2015.

Jovanovic, Boyan, "Firm-specific capital and turnover," *Journal of political economy*, 1979, *87* (6), 1246–1260.

_ , "Job matching and the theory of turnover," *Journal of political economy*, 1979, *87* (5, Part 1), 972–990.

Jr, Roland G Fryer, "Belief flipping in a dynamic model of statistical discrimination," *Journal of Public Economics*, 2007, *91* (5-6), 1151–1166.

Juhn, Chinhui and Kristin McCue, "Specialization then and now: Marriage, children, and the gender earnings gap across cohorts," *Journal of Economic Perspectives*, 2017, *31* (1), 183–204.

Kahn, Lisa B, "The long-term labor market consequences of graduating from college in a bad economy," *Labour Economics*, 2010, *17* (2), 303–316.

_ **and Fabian Lange**, "Employer learning, productivity, and the earnings distribution: Evidence from performance measures," *The Review of Economic Studies*, 2014, *81* (4), 1575–1613.

Kessler, Judd B., Corinne Low, and Colin D. Sullivan, "An Audit Alternative: Measuring Employer Preferences and Beliefs without Deception," *Working Paper*, 2018.

Khanna, Shantanu, "Salary History Bans and Wage Bargaining: Experimental Evidence," *Labour Economics*, 2020, p. 101853.

Kline, Patrick and Christopher Walters, "Audits as Evidence: Experiments, Ensembles, and Enforcement," *arXiv preprint arXiv:1907.06622*, 2019.

Lange, Fabian, "The speed of employer learning," *Journal of Labor Economics*, 2007, *25* (1), 1–35.

Laschever, Sara and Linda Babcock, *Women Don't Ask: Negotiation and the Gender Divide*, Princeton University Press, 2003.

Lazear, Edward P, "Firm-specific human capital: A skill-weights approach," *Journal of political economy*, 2009, *117* (5), 914–940.

Levy, Frank and Richard J Murnane, "With what skills are computers a complement?," *The American Economic Review*, 1996, *86* (2), 258–262.

Makridis, Christos and Maury Gittleman, "On the Cyclicality of Real Wages and Employment: New Evidence and Stylized Facts from Performance Pay and Fixed Wage Jobs," *Available at SSRN 3017034*, 2020.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2003.

Marianne, Bertrand, "New perspectives on gender," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1543–1590.

Mask, Joshua, "Salary History Bans and Healing Scars from Past Recessions," 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.

Mazei, Jens, Joachim Hüffmeier, Philipp Alexander Freund, Alice F Stuhlmacher, Lena Bilke, and Guido Hertel, "A meta-analysis on gender differences in negotiation outcomes and their moderators.," *Psychological bulletin*, 2015, *141* (1), 85.

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.

Moss-Racusin, Corinne A, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman, "Science faculty's subtle gender biases favor male students," *Proceedings of the National Academy of Sciences*, 2012, *109* (41), 16474–16479.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer, "Coarse thinking and persuasion," *The Quarterly journal of economics*, 2008, *123* (2), 577–619.

Murciano-Goroff, Raviv, "Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers," *Working Paper*, 2017.

Niederle, Muriel, "Gender. Handbook of Experimental Economics," 2015.

_ **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, "The making of an investment banker: Macroeconomic shocks, career choice, and lifetime income," *The Journal of Finance*, 2008.

_ , Scott Schaefer et al., "Personnel Economics: Hiring and Incentives," *Handbook of Labor Economics*, 2011, *4*, 1769–1823.

PayScale, "The Salary History Question: Alternatives for Recruiters and Hiring Managers," *White Paper*, 2017.

Ponthieux, Sophie and Dominique Meurs, "Gender inequality," in "Handbook of income distribution," Vol. 2, Elsevier, 2015, pp. 981–1146.

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium wage dispersion with worker and employer heterogeneity," *Econometrica*, 2002, *70* (6), 2295–2350.

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.

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

Rozada, Martín González and Eduardo Levy Yeyati, "Do women ask for lower salaries? The supply side of the gender pay gap," *Working Paper*, 2018.

Sherman, Eliot, Raina Brands, and Gillian Ku, "The Effect of Salary History Bans on Gender-Based Disparities in Compensation: Evidence From a Field Experiment," *Working Paper*, 2019.

Sinha, Sourav, "Salary History Ban: Gender Pay Gap and Spillover Effects," *Available at SSRN 3458194*, 2019.

Solnick, Sara J, "Gender differences in the ultimatum game," *Economic Inquiry*, 2001, *39* (2), 189–200.

Spence, Michael, "Job Market Signaling," *The Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

Sran, Gurpal, Felix Vetter, and Matthew Walsh, "Employer Responses to Pay History Inquiry Bans: Evidence from Online Job Postings, Hiring, and Pay," *Working Paper*, 2020.

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

Takahashi, Ana Maria and Shingo Takahashi, "Gender salary differences in economics departments in Japan," *Economics of Education Review*, 2011, *30* (6), 1306–1319.

Thilmany, Dawn, "Gender based differences of performance and pay among agricultural economics faculty," *Review of Agricultural Economics*, 2000, *22* (1), 23–33.

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

Westin, Alan F, "Past and future in employment testing: A socio-political overview," *U. Chi. Legal F.*, 1988, p. 93.

Wolinsky, Asher, "Prices as signals of product quality," *The review of economic studies*, 1983, *50* (4), 647–658.

Wozniak, Abigail, "Discrimination and the effects of drug testing on black employment," *Review of Economics and Statistics*, 2015, *97* (3), 548–566.

Figure 1: Ban Compliance Types

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

Figure 2: Callbacks and Amounts Disclosed

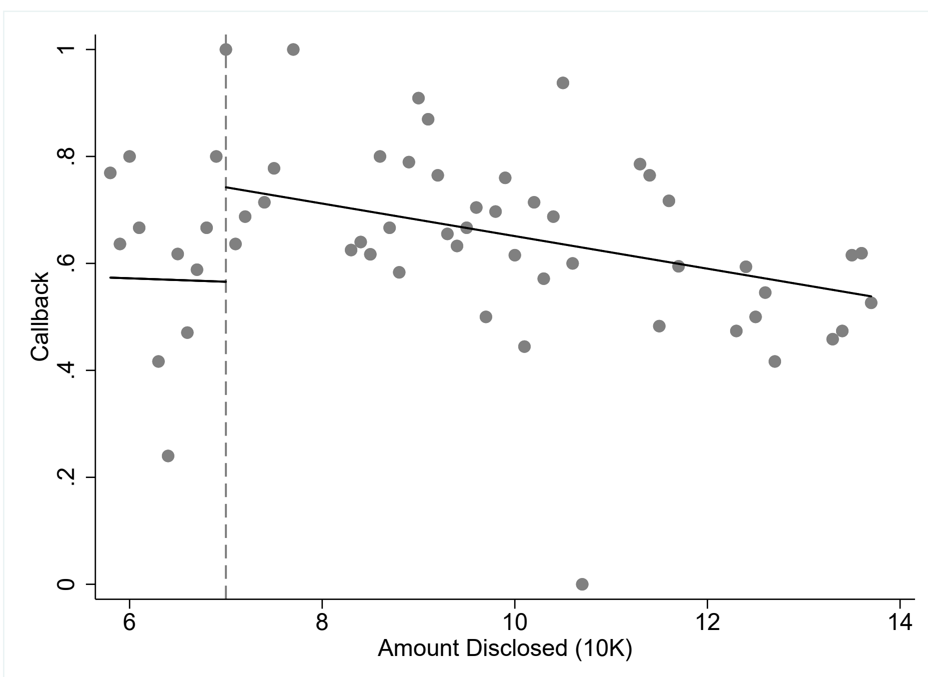

Notes: In this figure the circles present the proportion of job candidates who were recommended for a callback by the amount they disclosed. There is a line of best fit for observations in the lower 15% of the data, and one for the upper 85% of the data.

Table 1: Candidate Summary Statistics

	All			Male Candidate			Female Candidate		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	All	Salary Disclosed	Not Disclosed	All	Salary Disclosed	Not Disclosed	All	Disclosed	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) x Surplus	3,268	3,200	3,367	3,133	2,881	3,500	3,403	3,519	3,234
$\geq$ 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. Willingness-to-Pay (WTP) is the maximum a recruiter stated we should pay a particular candidate. Outside offer X %ile is their answer to the question "what offer would this candidate accept with X probability". Outside offer range is Outside Offer 95th %ile - Outside Offer 5th %ile. Offer is the take-it-or-leave-it-offer the recruiter suggested. Callback is whether the recruiter suggested the candidate be interviewed. Surplus is WTP minus Offer. p(accept) is the probability that a candidate would accept the specific take-it-or-leave it offer as approximated by fitting a logistic function through the X %ile answers. p(accept) x Surplus is the expected surplus from a certain take-it-of-leave it offer. $\geq$ 2 Other Offers is whether the recruiter believes this candidate would have 2 or more other offers (as opposed to 1 or 0). Offer | CB is the amount of the offer conditional on the recruiter recommending the candidate for a callback.

Table 2: Average Effect of Disclosing Salary

	(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	0.68***	0.97***	0.84***	1.30***	0.73***	1.08***	-0.92***	-1.25**
	(0.14)	(0.16)	(0.14)	(0.15)	(0.14)	(0.16)	(0.24)	(0.40)
Female x Disclosed		-0.57***		-0.92***		-0.69***		0.67+
		(0.13)		(0.11)		(0.12)		(0.37)
Female Disclosure Effect:								
Total		0.40		0.38		0.39		-0.59
p-value		0.01		0.01		0.01		0.00
Mean Non-Dislosers:								
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	0.18	0.18	0.25	0.26	0.20	0.21	0.01	0.02
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 5. 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 3: Average Effect of Disclosing a High versus Low Salary

	(1) WTP	(2) Outside Offer 50th %ile	(3) Offer	(4) Outside Offer Range
Disclosed 25th %ile Salary	-0.08	-0.04	-0.04	-0.52
	(0.14)	(0.13)	(0.12)	(0.34)
Disclosed 75th %ile Salary	1.11***	1.24***	1.18***	-0.37
	(0.15)	(0.14)	(0.13)	(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: 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 4: Average Effect of Disclosing by Salary Amount

	(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
Disclosed x Amt Disclosed	0.65***	0.67***	0.77***	0.78***	0.68***	0.69***	-0.03	0.03
	(0.06)	(0.06)	(0.05)	(0.05)	(0.05)	(0.06)	(0.12)	(0.18)
Salary Disclosed	-5.76***	-6.17***	-6.82***	-7.11***	-6.04***	-6.33***	-0.19	-1.12
	(0.56)	(0.61)	(0.49)	(0.54)	(0.53)	(0.56)	(0.92)	(1.68)
Female x Disclosed x Amt Disclosed		0.01		-0.00		0.02		-0.03
		(0.05)		(0.04)		(0.04)		(0.14)
Female x Disclosed		0.34		0.27		0.19		1.02
		(0.47)		(0.39)		(0.41)		(1.49)
Female Amount Disclosed Slope:								
Total		0.68		0.78		0.71		-0.00
p-value		0.00		0.00		0.00		0.99
Mean Non-Dislosers:								
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	0.38	0.38	0.54	0.54	0.44	0.44	0.01	0.01
Observations	2048	2048	2048	2048	2048	2048	2048	2048

Notes: 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 5 that include interactions with gender. 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. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table 5: Effect of An Extra Dollar Decomposed

	(1) WTP	(2) Outside Offer 50th %ile	(3) Offer	(4) O.O. Range	(5) p WTP = Outside Offer
+$10k from Firm	0.64***	0.77***	0.68***	-0.01	0.00
	(0.07)	(0.06)	(0.06)	(0.20)	
+$10k from Male	0.42***	0.62***	0.48***	-0.39	0.00
	(0.09)	(0.07)	(0.09)	(0.26)	
+$10k within Firm	0.70***	0.79***	0.73***	0.04	0.06
	(0.06)	(0.05)	(0.05)	(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: 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 6 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 1. 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, 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. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table 6: Average Effect of Disclosing for Callback Outcomes

	(1) Callback	(2) Callback	(3) Surplus	(4) Surplus	(5) p(accept)	(6) p(accept)	(7) p(accept) *x* Surplus	(8) p(accept) *x* Surplus
Salary Disclosed	-0.01	-0.04	-0.05	-0.11*	-0.01	-0.04	-0.01	-0.06+
	(0.02)	(0.03)	(0.04)	(0.04)	(0.02)	(0.03)	(0.03)	(0.03)
Female x Disclosed		0.07+		0.12**		0.06**		0.09**
		(0.04)		(0.04)		(0.02)		(0.03)
Female Disclosure Effect:								
Total		0.02		0.01		0.01		0.03
p-value		0.48		0.80		0.57		0.30
Mean Non-Dislosers:								
All	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
Male	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
Female	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 5. Dependent variables are listed in the column header and explained in notes to Table 1. Outcomes measured in dollars (e.g. surplus) are in $10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

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

	(1) Callback	(2) Callback	(3) Surplus	(4) Surplus	(5) p(accept)	(6) p(accept)	(7) p(accept) *x* Surplus	(8) p(accept) *x* Surplus
Salary Disclosed	-0.00	-0.03	-0.08*	-0.12*	-0.02	-0.04	-0.04	-0.06*
	(0.03)	(0.03)	(0.04)	(0.05)	(0.02)	(0.03)	(0.03)	(0.03)
Female x Disclosed		0.07		0.09*		0.05+		0.06+
		(0.04)		(0.04)		(0.02)		(0.03)
Female Disclosure Effect:								
Total		0.03		-0.03		0.00		-0.01
p-value		0.32		0.42		0.92		0.83
Mean Non-Dislosers:								
All	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
Male	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
Female	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 5 and mimics Table 6; 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. Outcomes measured in dollars (e.g. Surplus) are in $10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

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

	(1) Callback	(2) Surplus	(3) p(accept)	(4) p(accept) x Surplus
+$10k from Firm	-0.02	-0.04+	-0.02	-0.04**
	(0.02)	(0.02)	(0.01)	(0.01)
+$10k from Male	-0.05+	-0.07*	-0.03	-0.05*
	(0.03)	(0.03)	(0.02)	(0.02)
+$10k within Firm	0.00	-0.03	-0.01	-0.03+
	(0.02)	(0.02)	(0.01)	(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: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 6 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 1. 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". Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table 9: Average Effect of Disclosing by Salary Amount for Callback Outcomes

	(1) Callback	(2) Callback	(3) Surplus	(4) Surplus	(5) p(accept)	(6) p(accept)	(7) p(accept) *x* Surplus	(8) p(accept) *x* Surplus
Salary Disclosed	0.13	0.18	0.28	0.17	0.15	0.12	0.33*	0.26+
	(0.14)	(0.18)	(0.17)	(0.21)	(0.12)	(0.14)	(0.13)	(0.15)
Female x Disclosed		-0.17		0.15		0.03		0.10
		(0.17)		(0.19)		(0.10)		(0.12)
Disclosed x Amt Disclosed	-0.01	-0.02	-0.03*	-0.03	-0.02	-0.02	-0.03**	-0.03*
	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Female x Disclosed x Amt Disclosed		0.02		-0.01		-0.00		-0.01
		(0.02)		(0.02)		(0.01)		(0.01)
Female Amount Disclosed Slope:								
Total		0.01		-0.04		-0.02		-0.04
p-value		0.71		0.07		0.23		0.01
Mean Non-Dislosers:								
All	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
Male	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
Female	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: 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 5. 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. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table 10: Average Effect of Disclosing by Salary Amount for Callback Outcomes Salary > $70,000

	(1) Callback	(2) Callback	(3) Surplus	(4) Surplus	(5) p(accept)	(6) p(accept)	(7) p(accept) x Surplus	(8) 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:								
Total		-0.01		-0.02		-0.03		-0.03
p-value		0.46		0.44		0.08		0.04
Mean Non-Dislosers:								
All	0.64	0.64	0.65	0.65	0.56	0.56	0.34	0.34
Male	0.66	0.66	0.69	0.69	0.56	0.56	0.35	0.35
Female	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: The sample in this table is restricted to those that did not disclose a salary and those that disclosed a salary above $70,000. All models include recruiter controls and candidate fixed effects. This table shows estimates from versions of Equation 5. Dependent variables are listed in the column header and explained in notes to Table 1. Salary amounts and outcomes measured in dollars (e.g. surplus) are in $10K increments. Robust standard errors are clustered at the recruiter level. + $p < 0.10$ * $p < 0.05$ ** $p < 0.010$ *** $p < 0.001$

Table 11: Complier Type Cells

Type of Cells			*Proportion of Cell that is an ...*				*Mean Disclosure Δ*	
Gender	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.

Table 12: Effect of Bans (Full or Partial) on Callbacks

Callbacks

	Women	Men	Ratio
No Ban	.64	.63	1
	(.023)	(.023)	(.045)
Ban	.64	.63	1
	(.029)	(.024)	(.05)
Ban-No Ban	-.0034	-.0027	-.00096
	(.037)	(.033)	(.067)
p-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.

Table 13: Full vs. Partial Ban: Effect on Annual Salaries, Conditional on Callbacks

Salary | Callbacks (Full Ban)

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

Salary | Callbacks (Partial Ban)

	Women	Men	Ratio
No Ban	101824.98	112220.58	0.91
	(1111.61)	(1315.39)	(0.01)
Ban	97966.75	106061.10	0.92
	(319.37)	(478.83)	(0.01)
Ban-No Ban	-3858.23	-6159.48	0.02
	(1156.58)	(1399.83)	(0.01)
p-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.