Negotiation, Employer Demand, and Gender Disparities

Ratchanon Chotiputsilp, Taeho Kim, Clémentine Van Effenterre*

This version: July 2026

Abstract

We study how candidates' negotiation behavior shapes employer decisions using a field experiment in which 493 professional recruiters evaluated more than 7,000 job applications and counteroffers. Negotiation updates employer beliefs about both outside options and productivity. Employers revise outside-option beliefs similarly for male and female candidates, and wage gains from intensive negotiation are similar across genders. However, employers revise productivity upward for men and not for women. As a result, more intensive negotiation increases recommendations for higher-level roles substantially more for men than for women, and female candidates face larger increases in rejection rates than comparable male candidates. Combining these effects, we estimate optimal counteroffers of 11–12 percent above the initial offer for women and 20–30 percent for men. Pay transparency attenuates gender differences in the returns to negotiation in higher-level role recommendations by constraining recruiter discretion.

JEL codes: J08, J16, J24 **Keywords**: *Negotiation*; *Gender*; *Experiment*; *Transparency*

*__Chotiputsilp__: Erasmus University Rotterdam, chotiputsilp@rsm.nl and Thammasat University, ratchanon@econ.tu.ac.th, __Kim__: University of Toronto, thk.kim@utoronto.ca,__Van Effenterre__: University of Toronto, c.vaneffenterre@utoronto.ca. We thank seminar participants at Canadian Labour Economics Forum (CLEF) Conference, Simon Fraser University, SOLE, Texas A&M University, UBC, University of Toronto, and Wharton People and Organizations Conference for helpful discussions and comments. Nathalia Wyss and Garrett Evans provided outstanding research assistance. This project benefited from the support of the RSF Foundation. IRB clearance was obtained from the University of Toronto (RIS Human Protocol Number 45563). The survey experiment was pre-registered on May 13, 2024, ID: AEARCTR-0012700, a pre-analysis plan was uploaded on the AEA RCT Registry website on May 13, 2024.

Introduction

Individual-level bargaining is widespread in labor markets, even in settings with formal pay structures (Hall and Krueger, 2012; Lachowska et al., 2022; Caldwell, Haegele and Heining, 2023). While standard search and bargaining models emphasize workers' outside options as determinants of wages (Manning, 2011; Caldwell and Danieli, 2024; Jäger et al., 2024), employers rarely observe candidates' outside options or productivity directly during hiring. Instead, they may use negotiation behavior to infer strong outside options, high productivity, or both. These inferences can shape wages, career placement, and hiring decisions, yet we know little about how employers interpret negotiation behavior or how these beliefs affect hiring outcomes.

Bargaining is also a key source of gender inequality in the labor market. Differences in pay-setting practices contribute to the gender wage gap (Card, Cardoso and Kline, 2016; Biasi and Sarsons, 2022; Palladino et al., 2025), and a large literature documents gender differences in negotiation behavior and outcomes. However, it remains unclear whether these gaps reflect differences in workers' behavior (Bowles, Babcock and Lai, 2007; Amanatullah and Morris, 2010; Roussille, 2020; Exley, Niederle and Vesterlund, 2020) or differences in how employers interpret identical negotiation behavior. If employers draw different inferences from the same negotiation behavior across groups, observationally equivalent actions can generate divergent outcomes across workers.

To study these questions, we combine a conceptual framework, a large-scale field experiment with professional recruiters, and a structural model. The conceptual framework formalizes how employers infer candidates' outside options and productivity from negotiation behavior. The field experiment, comprising more than 7,000 recruiter evaluations of job applications and counteroffers, identifies these inferences and tests whether they differ by candidate gender. We then exploit experimental variation to examine the contextual determinants of gender gaps. Finally, we estimate the model structurally to isolate how much of the gender gaps is attributable to the belief asymmetry itself, derive candidates' optimal counteroffers, and quantify how pay transparency changes the firm's decision-making.

We start by outlining a partial equilibrium model of employer learning from negotiation behavior. Employers use negotiation intensity to infer candidates' outside options and pro-

ductivity, and these beliefs determine wages, career placement, and selection decisions—such as whether to recommend a candidate for a higher-level role or to proceed with the candidacy. These beliefs need not affect all decisions in the same way. When wage-setting is subject to institutional constraints (e.g., pay bands, internal equity pressures, and legal scrutiny) and productivity is complementary to higher-level roles, employers respond to these beliefs along different margins. Higher perceived outside options raise wages, while placement decisions depend more directly on perceived productivity. When a high counteroffer raises perceived outside options without a corresponding increase in perceived productivity, the candidate becomes less attractive to hire and the firm may reject the candidacy. Group differences in outcomes therefore depend on both how beliefs are updated across groups and how sensitive each margin is to those beliefs.

Studying employers' beliefs about observationally equivalent workers who negotiate requires data that remain scarce (Cullen, 2024). We address this gap using a field experiment in which 493 professional recruiters and HR professionals were hired and paid to provide hiring recommendations. They evaluated candidates who had received job offers, submitted counteroffers, and were actively engaged in negotiation, without being informed that the applications were fictitious. Participants made recommendations about pay, work arrangements, whether candidates should be considered for a concurrently open higher-level role, and whether the firm should continue with the candidacy. This setting mirrors real-world recruitment outsourcing, where firms delegate hiring decisions to third-party recruiters. Our approach builds on Agan, Cowgill and Gee (2021) by using real recruiters making consequential decisions, but extends this framework to study how employers interpret negotiation behavior and update beliefs. The data consist of structured evaluation forms capturing these recommendations, along with candidate assessments and qualitative evaluations, as well as a follow-up survey that elicits recruiters' beliefs.

To examine negotiation dynamics, we introduced four key levels of randomization. First, in each batch of 16 resumes that a recruiter evaluated, half of the candidates were assigned a female first name, and the other half a male first name. Second, within each batch, a proportion of randomly-selected candidates requested counteroffers above the target salary range, while the remaining candidates requested counteroffers below the range. Third, within each recruiter's resume packet, the proportion of candidates presenting counteroffers above

the maximum salary was either 25 percent, 50 percent, or 75 percent. Fourth, we use a cross-recruiter randomization to evaluate the impact of pay transparency on recruiters' decisions.

Our first set of results shows that counteroffers above the target range lead recruiters to revise beliefs about both candidates' outside options and productivity. Recruiters revise upward their estimates of current salaries and outside offers, and increase productivity-related assessments, including technical skills and initiative. Consistent with this, both salary recommendations and senior-track recommendations increase. At the same time, higher outside-option assessments lead recruiters to view intensively negotiating candidates as less likely to accept the offer, and increase the likelihood of recommending rejection.

Our second set of results reveals a sharp gender divergence. Men and women receive similar salary gains from more intense negotiation, consistent with nearly symmetric updating of outside-option beliefs. However, the returns to negotiation differ substantially for advancement and selection decisions. Negotiating above the target range increases senior-track recommendations for both genders, but significantly more for men, generating a gender gap of 5.9 percentage points. Female candidates also experience larger increases in rejection rates than comparable male candidates. Together, these findings suggest that focusing solely on wage outcomes understates the role of negotiation in generating gender inequality.

We next examine the contextual factors that shape these gender differences in the returns to intense negotiation. We document substantial heterogeneity across recruiters: the rejection penalty for female candidates with high counteroffers is concentrated among recruiters who exhibit more male-favoring evaluation patterns. We also find that gender differences in senior-track recommendations are driven by recruiters located in regions with higher implicit gender bias.

We then examine whether policy interventions can mitigate these gender differences. Pay transparency policies have expanded rapidly in recent years, motivated in part by the goal of reducing employer discretion and narrowing labor market disparities. We experimentally vary whether recruiters are informed that the salary range was disclosed to candidates. We find that transparency increases recruiters' awareness of the costs of offering salaries above the target range. While we find limited effects on salary levels during the negotiation stage overall, transparency meaningfully reduces gender differences in senior-track recommendations, consistent with limiting discretion in this part of the hiring process.

Finally, we build and estimate a structural model of the firm's hiring decision to move beyond these reduced-form patterns and quantify the sources of the observed gender asymmetry. The model allows us to hold recruiters' beliefs about a candidate's skill and outside option fixed and ask how much of each gender gap survives; equalizing beliefs across gender substantially narrows the gap in every condition we study, indicating that the asymmetry in perceived skill—recruiters revising upward their assessments of technical skills, team skills, and initiative for men who negotiate intensively, but not for women—is an important channel generating the gap, consistent with recent evidence that evaluators draw different inferences from identical signals depending on the sender's gender (Sarsons, 2019). We also use the structural estimates to characterize the counteroffer that maximizes each candidate's certainty equivalent, accounting for the full distribution of recruiter responses: the optimal counteroffer is 11–12 percent above the initial offer for women and 20–30 percent for men, a gap that arises because aggressive negotiation triggers belief updating for men but not for women. Re-estimating the model separately under pay secrecy and transparency then lets us characterize the mechanism behind the reduced-form transparency result: the firm's overall resistance to senior-track recommendations rises for men and women alike even as recruiters' beliefs remain unchanged—indicating that transparency operates by constraining firm discretion, rather than by changing how recruiters' beliefs translate into recommendations.

Our paper contributes to the literature on models of wage setting and bargaining (Burdett and Mortensen, 1998; Postel-Vinay and Robin, 2002; Cahuc, Postel-Vinay and Robin, 2006; Alan, 2011; Hall and Krueger, 2012; Jäger et al., 2020) and their implications for pay differentials (Brenzel, Gartner and Schnabel, 2014; Card, Cardoso and Kline, 2016; Harmon and Caldwell, 2019; Lachowska et al., 2022; Card, 2022; Morchio and Moser, 2023; Caldwell, Haegele and Heining, 2023). Existing bargaining models emphasize how the parties' outside options determine negotiated wages, while empirical work has examined how bargaining contributes to wage inequality. In these frameworks, productivity and outside options are largely treated as inputs into the bargaining process. Our results show that bargaining also shapes employers' beliefs about these inputs. Holding candidate qualifications constant while varying negotiation behavior, we find that a counteroffer updates employers' perceptions of both candidates' outside options and productivity, and that these belief updates govern different employer decisions. Consistent with canonical bargaining models, wages respond primarily

to perceived outside options, whereas senior-track recommendations respond more directly to perceived productivity. Negotiation therefore affects not only how match surplus is divided through wages, but also how employers assess candidates' long-run value to the firm and, consequently, their career trajectories.

This paper also contributes to the literature on gender differences in bargaining (Niederle and Vesterlund, 2007; Babcock et al., 2013; Hall and Krueger, 2012; Leibbrandt and List, 2015; Card, Cardoso and Kline, 2016; Säve-Söderbergh, 2019; Roussille, 2020; Exley, Niederle and Vesterlund, 2020; Biasi and Sarsons, 2022; Cullen, Pakzad-Hurson and Perez-Truglia, 2025) by unpacking employers' decision-making during the negotiation stage of hiring. Among recent studies examining negotiation in a labor market setting, Roussille (2020) focuses on initial salary requests, while recent RCTs on negotiation training (Chotiputsilp and Kim, 2023; Cullen, Pakzad-Hurson and Perez-Truglia, 2025) assess overall candidate-firm interactions. In contrast, we isolate the phase in which recruiters respond to counteroffers, after the screening process (e.g., interviews, test results) is complete. Focusing on counteroffers allows us to capture a more direct measure of negotiation, as this stage reflects active employer-candidate bargaining rather than initial signaling through salary expectations. The back-and-forth negotiation phase provides a critical window for employers to differentially interpret assertiveness and to make compensation decisions (Caldwell, Haegele and Heining, 2023). By combining this targeted setting, variation in negotiation intensity, and detailed recruiter assessments, we are able to analyze how different types of information are conveyed and interpreted during bargaining. This approach offers a novel and rich view into the mechanisms driving gender differences in negotiation outcomes.

Our experiment uncovers a new mechanism through which gender discrimination can arise in hiring: recruiters' assessments of candidates' skills—not just their hiring or pay decisions—respond differently to identical negotiation behavior depending on candidate gender, and this differential interpretation of the same signal generates divergent advancement decisions later in the hiring process. This complements a literature that has found mixed evidence of gender discrimination in callback rates from large-scale audit studies (Bertrand and Duflo, 2017; Kline and Walters, 2021; Kline, Rose and Walters, 2022*a*). The absence of gender differences at the initial callback stage need not imply the absence of discrimination if the relevant decisions occur later, in how employers evaluate candidates' productivity and advancement potential.

Our design provides unusually direct evidence on statistical discrimination (Phelps, 1972; Arrow, 1973) by allowing us to observe both the signal (the counteroffer) and employers' resulting belief revision. Combining these data with a structural model, we quantify how much of the downstream gender gap this channel can explain. We further document substantial heterogeneity across recruiters, highlighting discretion as an important source of inequality in labor market outcomes (Hoffman, Kahn and Li, 2018; Adler, 2020; Agan, Cowgill and Gee, 2021; Haegele, 2022; Benson, Li and Shue, 2024; Minni et al., 2026).

Finally, our paper contributes to documenting the role of policy in shaping negotiation outcomes. Research has indicated that reducing information friction regarding pay can decrease wage disparities (Mas, 2017; Burn and Kettler, 2019; Bennedsen et al., 2022; Seitz, Gulyas and Sinha, 2022; Duchini, Simion and Turrell, 2022; Baker et al., 2023; Cullen and Pakzad-Hurson, 2023). When firms are required to disclose expected compensation, they have less latitude to differentially set pay, potentially curbing inequities. Understanding how these policies influence negotiation dynamics is critical for designing more effective interventions. However, to date, no field data have captured candidate–employer bargaining under varying transparency regimes, so the mechanisms behind these effects remain unclear.

The rest of the paper is structured as follows. Section 1 presents the model. In Section 2, we discuss the experimental design and provide summary statistics. In Section 3, we present the econometric specifications. In Section 4, we present results on returns to negotiating more intensely and differential returns by gender and Section 5 examines how these differences are influenced by contextual factors such as recruiter characteristics and pay transparency regimes. Section 6 structurally estimates the model, decomposes the sources of gender gaps, and derives candidates' optimal counteroffers. Section 7 concludes.

1 Conceptual Framework

We build a partial equilibrium model that motivates the empirical analysis in Section 2 and that we structurally estimate in Section 6. The model has three key ingredients: (1) negotiation conveys information about both skill and outside options; (2) firms respond to updated beliefs through multiple hiring decisions—particularly wages and senior-track recommendations—which may respond differently due to constraints on wage-setting and the greater role of

productivity in senior-track roles; (3) employers may interpret negotiation differently across demographic groups. Appendix A provides details and proofs.

1.1 Environment

Candidate attributes. A candidate belongs to group $g \in \{m, f\}$ and has latent skill v and outside option ω. The joint prior over (v,ω) is group-specific and Gaussian:

$$\begin{pmatrix} v \\ \omega \end{pmatrix} \Bigg| g \sim \mathcal{N}\left(\begin{pmatrix} \mu_{v,g} \\ \mu_{\omega,g} \end{pmatrix}, \Sigma_g \right), \qquad \Sigma_g = \begin{pmatrix} \sigma^2_{v,g} & \sigma_{v\omega,g} \\ \sigma_{v\omega,g} & \sigma^2_{\omega,g} \end{pmatrix}. \tag{1}$$

We assume that $\sigma_{v\omega,g} \neq 0$ which allows skill and outside options to be correlated within each group (e.g., higher-skill candidates may systematically have better outside alternatives).

Negotiation as a signal. The employer observes negotiation intensity n generated by

$$n = \beta_{v,g}\, v + \beta_{\omega,g}\, \omega + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma^2_\varepsilon), \tag{2}$$

where $\beta_{v,g} \geq 0$ measures how informative negotiation is about skill for group g, and $\beta_{\omega,g} \geq 0$ measures how informative it is about outside options. This structure allows negotiation to be differentially informative about both skill and outside options across groups.

Given (n,g), the employer forms a posterior distribution over (v,ω), denoted $f_g(v,\omega \mid n)$. Under the normal-linear signal structure and Gaussian priors, the posterior is jointly normal (see Appendix A.1).

Candidate's acceptance decision. The candidate accepts with probability

$$P(\text{accept} \mid w, s, \omega) = \Lambda(w + \tau s - \omega), \qquad \Lambda(x) = \frac{1}{1 + e^{-x}}, \tag{3}$$

where $\tau \geq 0$ is the candidate's utility value of senior-track placement.

Organizational constraints. Organizations rarely have complete flexibility in determining compensation or role assignment. Instead, they operate within pre-approved pay bands and hiring plans that specify benchmark compensation and a target seniority level for the vacancy.

Deviating from these benchmarks requires organizational justification and may generate administrative, financial, and internal-equity costs.

For compensation, we define the relevant reference point as a convex combination of the institutional anchor and the candidate's inferred outside option:

$$w_g^{\text{eff}}(n) \equiv \alpha\,\hat{\omega}_g(n) + (1-\alpha)\,\bar{w}^{\text{inst}}, \qquad \alpha \in [0,1], \tag{4}$$

where $\hat{\omega}_g(n)$ is the posterior mean outside option defined in (A.4) and $\bar{w}^{\text{inst}}$ is the institutional pay-band anchor. The parameter α captures the extent to which the firm's internal reference point is pulled toward the candidate's inferred market value — reflecting that compensation decisions are benchmarked not only against internal equity but also against what the firm believes the candidate can obtain elsewhere. Deviations from w_g^{eff} may require additional budgetary justification, create internal equity concerns, or conflict with established compensation policies.

For role assignment, senior-track placement entails organizational costs that increase with the extent of the departure from standard placement. Such recommendations may require additional coordination with hiring managers and human resources, revisions to job classifications or headcount plans, and departures from the organization's intended staffing plan. We capture this through a quadratic cost in s, so that larger deviations from the default placement incur disproportionately higher justification and coordination costs.

We capture these organizational constraints in the following way:

$$C_w(w) = \frac{\kappa_w}{2}\Big(w - w_g^{\text{eff}}\Big)^2, \qquad C_s(s) = \frac{\kappa_s}{2}s^2, \qquad \kappa_w, \kappa_s > 0. \tag{5}$$

The parameters κ_w and κ_s measure the rigidity of compensation and role assignment, respectively. Larger values correspond to organizations with stricter pay-band policies or more rigid hiring and promotion procedures, making departures from institutional benchmarks increasingly costly.

Rejection and production. If the candidate rejects, the firm incurs a cost $c_R \geq 0$ (vacancy costs, recruiter penalties). If hired, output is $v(1+\delta s)$, where $\delta \geq 0$ captures productivity complementarity between skill and senior-track roles.

1.2 Firm's Optimization Problem and Optimal Response

The firm chooses (w, s) to maximize expected profit over the posterior beliefs $f_g(v, \omega \mid n)$:

$$\Pi(w, s \mid n, g) = \int\int \Big[\Lambda(w + \tau s - \omega)\, S(v, w, s) - \Big(1 - \Lambda(w + \tau s - \omega)\Big) c_R\Big]\, f_g(v, \omega \mid n)\, dv\, d\omega, \quad (6)$$

where $S(v, w, s) = v(1 + \delta s) - w - C_w(w) - C_s(s)$ is the net surplus from acceptance.

The firm makes an offer only if $\Pi(w^*, s^* \mid n, g) \geq \bar{V}$; otherwise, it rejects the candidacy. Appendix A.3 derives the first-order conditions formally.

Comparative statics: How beliefs affect outcomes. Differentiating the first-order conditions with respect to posterior beliefs yields:

Wages. Both perceived productivity $\hat{v}$ and perceived outside options $\hat{\omega}$ increase wages. Higher skill raises surplus; higher outside options require compensation to secure acceptance. Wage rigidity κ_w dampens both responses.

Senior-track recommendations. Perceived productivity $\hat{v}$ unambiguously increases recommendations: higher skill raises the marginal productivity return δv from senior placement. The effect of perceived outside options $\hat{\omega}$ is ambiguous. Higher outside options increase the value of senior placement for securing acceptance, but they also reduce the expected productivity gain—since high outside-option candidates are less likely to accept and realize that gain.

Differential instrument responses. Wages and senior-track recommendations may respond differently to belief updates. Specifically, under some plausible conditions —candidates value advancement ($\tau > 0$), productivity is complementary to senior roles ($\delta > 0$), wage-setting faces tighter institutional rigidity than senior-track placement, and acceptance rates are high—senior-track recommendations respond more strongly to perceived productivity than wages do (Appendix A.4): $\frac{\partial s^*}{\partial \hat{v}} > \frac{\partial w^*}{\partial \hat{v}}$.

These conditions are plausible in professional labor markets: senior roles offer complementary productivity gains, candidates value career progression, and wage-setting faces institutional constraints (e.g., pay bands, internal equity pressures, and legal scrutiny).

For outside options, wages respond directly through the acceptance channel, while the effect on senior-track recommendations is ambiguous due to opposing forces.

1.3 Group Differences: From Beliefs to Outcomes

How does negotiation affect outcomes differently by gender? For any outcome $y \in \{w^*, s^*\}$, the total response to negotiation intensity is:

$$\frac{dy_g^*}{dn} = \frac{\partial y^*}{\partial \hat{\omega}} \frac{d\hat{\omega}_g}{dn} + \frac{\partial y^*}{\partial \hat{v}} \frac{d\hat{v}_g}{dn}.$$

Subtracting across groups reveals two sources of gender gaps:

$$\frac{dy_m^*}{dn} - \frac{dy_f^*}{dn} = \frac{\partial y^*}{\partial \hat{\omega}} \underbrace{\left(\frac{d\hat{\omega}_m}{dn} - \frac{d\hat{\omega}_f}{dn} \right)}_{\text{outside-option updating gap } (\Delta_\omega)} + \frac{\partial y^*}{\partial \hat{v}} \underbrace{\left(\frac{d\hat{v}_m}{dn} - \frac{d\hat{v}_f}{dn} \right)}_{\text{productivity updating gap } (\Delta_v)} . \tag{7}$$

The first term captures differential updating of outside-option beliefs (arising from $\beta_{\omega,m} \neq \beta_{\omega,f}$), while the second captures differential productivity updating ($\beta_{v,m} \neq \beta_{v,f}$).

1.4 Testable Implications

The model delivers three sets of testable implications linking negotiation intensity, belief updating, and employer decisions.

Implication 1: Effects of negotiation on beliefs and decisions. More intensive negotiation raises both outside-option and productivity assessments. Because these beliefs affect decisions differently, outcomes need not move together: wages increase with higher perceived outside options, while senior-track recommendations depend more directly on perceived productivity. That is,

$$\frac{\partial w^*}{\partial \hat{\omega}} > 0, \qquad \frac{\partial s^*}{\partial \hat{v}} > 0.$$

The effect of outside-option beliefs on senior-track recommendations is ambiguous. When a high counteroffer raises perceived outside options without a corresponding increase in perceived productivity, expected surplus falls and rejection becomes more likely.

Implication 2: Belief updating and group differences. Equation (7) shows that differences in outcomes across groups arise from differences in belief updating. For outcome $y \in \{w^*, s^*\}$,

the gender gap in the response to negotiation decomposes as

$$\frac{dy_m^*}{dn} - \frac{dy_f^*}{dn} = \frac{\partial y^*}{\partial \hat{\omega}} \Delta_\omega + \frac{\partial y^*}{\partial \hat{v}} \Delta_v.$$

When outside-option updating differs across groups, wage responses should also differ. When productivity updating differs, gaps emerge more clearly in senior-track recommendations, which depend more directly on productivity beliefs. Differences in rejection rates arise when productivity updating diverges, which affects expected surplus.

Implication 3: Pay transparency and institutional rigidity. Pay transparency increases the visibility of deviations from salary benchmarks, raising scrutiny over employer decisions. When recruiters are informed that salary ranges were disclosed, constraints become tighter. As a result, transparency should dampen differential responses when beliefs are updated differently across groups. Formally, $|\partial y^*/\partial \hat{v}|$ and $|\partial y^*/\partial \hat{\omega}|$ are both decreasing in institutional rigidity κ (Appendix A.4).

2 Experimental Design

Our experimental design mimics a real-world recruitment process where a firm, in the midst of a large-scale ramp-up of hiring, outsources part of their HR recruiting process. ITo ensure realism, we consulted with three professional recruiters and collaborated closely with one participant from our study pool during the design phase. This consultant also reviewed all study materials to ensure that the language and evaluation tasks aligned with industry recruiting practices.

A trade-off in creating a realistic setting is the use of deception. The design shares features with resume audit studies, and thus relates to recent critiques of such approaches. One concern is that participants may learn from study materials in ways that generate spillovers to real-world hiring decisions. We considered an incentivized resume-rating design, such as Kessler, Low and Sullivan (2019), but this approach is difficult to implement in a setting involving negotiation without introducing experimenter demand effects. Our setting mitigates these concerns in two ways. First, we recruit external professional evaluators rather than embedding the study within firms' actual hiring processes, so the evaluations have no consequences for

real hiring decisions. Second, unlike traditional audit studies, our design does not involve real candidate interactions or realized outcomes. Participants do not observe whether candidates accept offers or how they perform, limiting the potential for learning that could affect future hiring behavior.

2.1 Recruitment and Incentives

Between May 2024 and November 2024, we invited all recruiters and HR professionals from a popular online platform for employers to hire contractors, located in the U.S. and Canada with hourly rates below $150. This threshold encompasses the large majority of US-based recruiters on the platform (approximately 90 percent of whom charge below $100 per hour). We ultimately collected our data from 493 recruiters and HR professionals.

In our recruitment process (see Appendix D for details), we aimed to mimic how companies typically approach recruiters when outsourcing portions of their hiring. Over the past two decades, this practice has become increasingly common as companies have delegated more of their recruiting efforts. Firms hire individual recruiters on a temporary basis or outsource the entire recruiting function to a third-party organization (Landay and DeArmond, 2018; Black, Hasan and Koning, 2020; Cappelli, 2019; Cowgill et al., 2021). According to Cappelli (2019), about 40 percent of firms have outsourced much if not all of their hiring process to recruitment process outsourcing firms (RPO), which in turn rely on subcontractors.

These outsourcing arrangements take several forms. End-to-end RPO transfers the entire hiring function to an external provider, while project-based RPO engages contractors for a defined hiring campaign. Modular RPO—the most relevant form for our purposes—outsources specific components of the recruitment process, such as candidate sourcing, screening, or the provision of market intelligence on compensation trends and competitor hiring patterns. This form is particularly common among firms with existing talent acquisition teams that require supplemental capacity in targeted areas (PeopleScout, 2023; Robert Walters, 2023).

In such settings, as Agan, Cowgill and Gee (2021) document, when firms rely on external recruiters to evaluate candidates, contractors are typically asked to screen applications and provide compensation recommendations, making the offer and negotiation stage a natural component of the contracted scope in settings such as ours.

Turning to the characteristics of our sample, Table 1 presents summary statistics of the

sample of recruiters. The majority of recruiters in our sample are female (74 percent) and white (53 percent). Most have completed at least a bachelor's degree (89 percent) and have more than five years of professional experience (77 percent). Study participants were paid their hourly wages posted on their profile pages. The mean hourly rate paid, $64, sits between the Bureau of Labor Statistics' 2023 mean wages for human resources specialists ($37) and human resources managers ($74), suggesting our sample is broadly representative of HR workers in terms of compensation.

Each participant was informed that they would evaluate 16 candidates as a consultant in the later stages of recruitment. We explained that a rapidly growing technology firm needed help assessing software engineering applicants and finalizing compensation packages. Recruiters were instructed not to seek any information about candidates online, and all communication took place through the platform's built-in messaging system.

Recruiters received only their hourly wage, with no additional bonuses. We avoided outcome-based incentives (e.g., tied to candidate retention) because they could induce recruiters to evaluate candidates by perceived acceptance likelihood rather than candidate quality — consistent with standard practice on the platform, where recruiters are paid a fixed rate and are incentivized instead by reputation and repeat business. In Section 4.4.2, we show that HR professionals in our sample paid attention to the materials we provided. For instance, they respond to specific keywords in the resumes and to key details that are indicative of candidate quality, and they evaluate applicants accordingly.

Procedures After a potential recruiter accepted our job offer, each recruiter received four materials. First, they received a job offer prompt that outlined their instructions and provided details about the employer's profile as a software company. We indicated that informal offers with a salary of $130,000 had been extended to candidates and mentioned that some expressed a willingness to negotiate with suggested counteroffers. The description stated our aim to navigate these negotiations effectively and ensure competitive compensation packages tailored to each candidate's unique profile. It also highlighted that all candidates had successfully passed rigorous technical assessments. Recruiters were informed of the budgeted salary range for the position ($115,000 to $155,000) and advised that offers outside this range could be

considered for exceptional candidates.[1]

Second, we provided them with a document showing the job advertisement we stated we had posted. This document details the job responsibilities and outlines the profile of ideal candidates.

Third, they received a packet containing 16 resumes, each representing a candidate profile within the software development industry. Before the study began, we created 32 biographies separated into two packets to ensure that our study participants reviewed a diverse group of candidates. Each recruiter received one randomly selected packet of 16 resumes. These fictitious profiles were based on actual resumes obtained from popular job board sites. The resume characteristics were derived from 208 actual resumes of software engineers with 3-10 years of experience from major cities in the US, all of whom had education and work experience in the U.S. We ensured that the career progressions in the randomized biographies made sense.

Fourth, they received a list of counteroffers, which included the candidate ID, name, and the counteroffer amounts proposed by the fictitious candidates. To ensure recruiters could provide actionable input, each of the 16 candidates on the list was designed to have submitted a counteroffer, as candidates without counteroffers did not require further consideration.

2.2 Randomization

A key randomized feature in this study is the candidate's gender, which is assigned using either a female or male first name. To avoid confounding gender with racial perceptions, we used names that are generally perceived as White or racially neutral. In each recruiter's packet, a randomly selected half of the 16 biographies are assigned a female name, while the other half are assigned a male name.

Additionally, we introduced a sub-treatment of negotiation intensity. Attached to each candidate's resume, we provided counteroffer salaries to the firm's initial offer of $130,000. These counteroffers, all above the initial offers, were randomly assigned and represent negotiation intensity. A certain portion of the candidates were shown to present counteroffers below the maximum salary range (but above the initial offer), while others presented counteroffers above the maximum, up to 20 percent above the threshold (i.e., from $156,000 to $185,000). Within each recruiter's resume packet, the proportion of candidates presenting counteroffers

[1] Appendix D presents the job posting.

above the maximum salary was either 25 percent, 50 percent, or 75 percent. The distribution of counteroffers is presented in Figure B1.

To understand how pay transparency affects employer behavior, we divided the recruiters into two groups. Half of the recruiters were assigned to the transparency treatment, where they were informed that the employer's job advertisements displayed the target salary range. The remaining recruiters were part of the no-transparency treatment, where they received information about the target salaries without any indication of transparency in the job adverts.

Recruiters were prompted the following depending on their treatment condition:

Our allocated salary range for these roles is $115,000 to $155,000. However, we are open to considering offers outside this range for exceptional candidates.

[Treatment: Posted Range] *To adhere to local labor laws concerning pay transparency in certain offices, we have explicitly mentioned this range in our job advertisement.*

[Treatment: No Posted Range] *We did not mention the target salary range in the job advertisement.*

In addition, recruiters in the transparency treatment received job advertisements that purportedly included the salary range for the position, whereas those in the no-transparency treatment received job advertisements that included the phrase "Competitive salary based on experience and qualifications" instead.

Because recruiters participate sequentially, there is a risk that standard randomization may not achieve balance across recruiter covariates. To mitigate this, we implemented a re-randomization process similar to that used by Agan, Cowgill and Gee (2021). Each time a new batch is launched and random treatment statuses are assigned, we conduct a balance test. For instance, we performed this test on the sample of recruiters collected to date by regressing recruiters' characteristics (e.g., experience, race, and gender) on indicators for the randomized treatments (e.g., percentage of candidates requesting salaries above the stated range, biography packet version 1 or 2, and transparency treatment). If the balance check fails (i.e., the p-value is less than 0.2), we re-randomized the treatment status and experimental material until balance is achieved.

We perform a series of standard data quality checks to ensure that the randomization was successful and that the experimental design was not compromised. For example, we test

whether resumes with counteroffers exceeding the maximum salary range were more likely to be evaluated by certain types of recruiters or had different resume characteristics. Table B1 shows that the types of recruiters (e.g., experience, race, and gender) evaluating counteroffers made below or above the target are similar, and that resumes with higher counteroffers are just as likely to have female names as male names. Additionally, we perform a balance test by regressing recruiters' characteristics on indicators for the pay transparency treatment. The results, presented in Table B2, confirm that recruiter characteristics on the platform are balanced across the transparency treatment and the control group.

First Survey Recruiters completed a structured questionnaire designed to capture their perceptions of each candidate (see Appendix F.1). To maintain realism, we implemented the survey using Google Forms rather than platforms such as Qualtrics. Although less common in academic research, Google Forms more closely resembles tools used in practice and reduces the likelihood that participants perceive the exercise as an academic study.

The questionnaire begins with evaluations of qualitative characteristics such as technical skills, teamwork skills, and initiative, measured using a 4-point Likert scale to assess levels of agreement or disagreement. To capture beliefs about outside options, we asked for recruiters' opinions on the candidates' potential other offers and their best estimate of current salaries.

Recruiters then provided compensation recommendations, including proposed salary offers and work-from-home arrangements, and indicated whether they would recommend each candidate for a concurrently open senior position.[2] Finally, they reported how they would respond to each candidate's counteroffer.

Post-Experiment Survey Once recruiters completed the evaluation of each profile, we asked them fill out a post-experiment survey (see Appendix F.2). The survey is designed to measure recruiters' perceptions about the overall quality of the pool of candidates and their perceptions of the behavior exhibited by other employers in the market. We asked them about their perceived prevalence of negotiation, as well as potential reasons why recruiters might be reluctant to accommodate wage negotiations.

[2]Redirecting a strong candidate to a concurrently open, more senior requisition is a familiar margin of recruiter discretion: applicant-tracking systems support moving candidates across requisitions (Cornerstone OnDemand, 2024), and organizations commonly post companion requisitions spanning adjacent levels.

2.3 Summary Statistics

Table 1 presents summary statistics for the recruiters who participated in our study. Seventy-four percent of the recruiters are female, and half are White. On average, they have logged 366 hours on the platform, with an hourly rate of $64. Eighty-nine percent hold bachelor's degrees, and 77 percent have five or more years of recruiting experience.

The counteroffers made by candidates are distributed across recruiters in accordance with our expectations for correct randomization. At the recruiter level, the average proportion of resumes with counteroffers exceeding the salary range is approximately 50 percent. Additionally, the likelihood of receiving resume packets with 25 percent, 50 percent, or 75 percent of candidates requesting higher salaries than the salary range is similar across all three groups. Finally, half of the recruiters were assigned to the pay transparency treatment.

2.4 External Validity and Generalizability

The generalizability of our findings may be affected by concerns related to selection, attrition, naturalness, and scale (List, 2020). Regarding selection, our sample of recruiters is comparable to the sample of HR professionals in the Current Population Survey (CPS) in terms of gender composition and average compensation. Attrition in our study was minimal (1.7 percent) and does not vary systematically with recruiters' observable characteristics (see Table B3). With respect to external validity, our design uses a 50:50 male-to-female ratio within each batch of 16 candidates to ensure sufficient statistical power. This differs from the gender distribution in the software engineering workforce, where women represent approximately 27 percent of workers (Amer, Craig and Van Effenterre, 2024). However, this feature of the design is unlikely to bias our results upward. If anything, it may lead us to understate the extent of gender disparities. First, a higher share of female candidates may signal a stronger emphasis on diversity, which would tend to attenuate discriminatory behavior; any bias we detect should therefore be interpreted as a lower bound. Second, if recruiters interpret a higher female share as indicative of lower average quality among female candidates, this would be reflected in lower overall ratings. Instead, we later show that, on average, female candidates are, if anything, evaluated more favorably on observable skill dimensions, absent heterogeneity by negotiation intensity.

3 Econometric Specifications

3.1 The Impact of Negotiation Intensity on Labor Market Outcomes

To explore how firms' compensation decisions differ based on the intensity of negotiation, we use the following specification to test the difference between candidates requesting salaries above the maximum salary range and those requesting salaries below it:

$$Y_{ij} = \beta_3 \times OverTarget_{ij} + \beta_4 \times X_j + \delta_i + \epsilon_{ij} \tag{8}$$

where Y_{ij} is the outcome of job seeker i who is reviewed by recruiter j, and $OverTarget_{ij}$ is a dummy variable equal to one if the counteroffer exceeds the salary target. X_j includes recruiters' controls, such as the gender, race, experience level. δ_i denotes candidate biography fixed effects. In all of our regressions, we cluster standard errors at the recruiter level.

Our primary outcomes capture multiple dimensions of recruiter decision-making. These include the recommended salary offer, whether the recruiter recommends considering the candidate for a newly available higher-level role, whether the recruiter recommends allowing the candidate to work remotely, and whether the recruiter recommends rejecting the candidacy. We also examine recruiters' assessments of candidates' current salaries, outside offers, likelihood of accepting the recommended offer, and several productivity-related evaluations, including technical skills, team skills, and initiative. Unless otherwise noted, binary outcomes are coded as indicators for favorable evaluations. For example, the senior-role outcome equals one if the recruiter strongly or moderately recommends considering the candidate for a higher-level role and zero otherwise.

To test whether the effect of requesting a salary above the target range differs based on the candidate's gender, we look at the interaction with the gender of the candidate by including an interaction term between the negotiation intensity and gender:

$$\begin{aligned} Y_{ij} &= \beta_5 \times Female_{ij} \times OverTarget_{ij} + \beta_6 \times Female_{ij} \\ &+ \beta_7 \times OverTarget_{ij} + \beta_8 \times X_j + \delta_i + \epsilon_{ij} \end{aligned} \tag{9}$$

where and $Female_{ij}$ is a dummy variable equal to one if the gender of the candidate is

female.

Finally, we examine pre-registered heterogeneity in recruiter evaluations based on recruiter characteristics and the conditions under which they assess candidates. For example, to investigate how concentrated or diffused recruiters' propensities are, we first construct a male-favoring propensity score for each recruiter and analyze heterogeneity across recruiter types. We use a pre-registered method to assign male-favoring propensities to recruiters. From the post-experimental survey, we asked each recruiter to name the top three candidates. We then compile a list of the top 6 biographies mentioned by all recruiters and calculate for each recruiter the ratio of male resumes selected to the total number of top 6 resumes presented to them with male names. Figure B2 shows that some biographies are disproportionately mentioned as top choices. Among the 32 resumes, some were mentioned fewer than 20 times, while others were mentioned more than 80 times. The number of times each biography was presented to recruiters ranges from 231 to 252, with a mean of 242. A predictor for a resume being one of the top 6 is longer job experience. Since it is possible that a recruiter may never be shown a top male candidate by chance, we add one to the denominator to avoid division by zero. If a recruiter did not rank anyone, we treat this as missing data for their propensity. The distribution of male-favoring propensities across recruiters is shown in Figure B3. We then analyze heterogeneity by recruiter propensities, excluding any resumes based on the top 6 biographies.

We also examine heterogeneity based on the prevalence of negotiation within a batch seen by a recruiter. Each packet was randomized such that 25 percent, 50 percent, or 75 percent of candidates may ask for counteroffers above the target salary range, while the rest submitted counteroffers below the range. To analyze how this prevalence affects evaluations, we further examine heterogeneity by low prevalence (25 percent) and high prevalence (50 percent and 75 percent) groups.

3.2 Pay Transparency Policy

To test the impact of the pay transparency policy, we use the following specifications:

$$Y_{ij} = \beta_1 \times Transparency_{ij} + \beta_2 \times X_j + \delta_i + \epsilon_{ij} \tag{10}$$

where $Transparency_{ij}$ is an indicator variable equal to one if the recruiter was informed that the advertisement listed a posted salary range. We also test an interaction model with indicators from the main specification that examines gender differences.

4 Results: Negotiation Intensity and Gender Disparities

4.1 Descriptive Evidence on Negotiation

In the post-experimental survey, we asked recruiters a series of questions to explore their views on negotiation.[3] Table 1 presents novel evidence about wage setting from these responses. Ninety-one percent of recruiters indicated that counteroffers above the target salary range are common, highlighting the prevalence of intense negotiation. However, when asked specifically about female candidates, only 62 percent of recruiters reported that intense counteroffers are common for women. This suggests that recruiters perceive a lower willingness to negotiate among female candidates compared to male candidates.

The survey responses also support the two informational channels emphasized in the model. Sixty-two percent of recruiters report that counteroffers reveal information about current salaries and outside options, while 43 percent say they reveal something about the candidate's personality. Recruiters also report meaningful constraints in accommodating high requests: 66 percent cite concerns about broader renegotiation within the organization, and 65 percent cite fairness concerns across candidates. A smaller share (26 percent) mentioned legal consequences for making exceptions to the salary range.

These responses suggest that counteroffers convey information about candidates that may shape negotiation outcomes, while organizational constraints limit how firms can respond to that information.

4.2 Signaling Value of Higher Counteroffers

The framework in Section 1 implies that negotiation intensity updates employer beliefs about both outside options (ω) and skill (v), which in turn shape pay, senior-track recommendations,

[3]These questions were asked after participants completed the main survey to avoid any experimenter demand effects.

and whether recruiters choose to continue with the candidacy. We begin by examining these patterns in the overall sample.

Table 2 shows that counteroffers exceeding the indicated salary target convey higher market value. When a candidate asks above the announced target—an increase of roughly 20 percent, from \$142,000 to \$170,000—recruiters recommend significantly higher salary offers, increasing by \$14,567 (10.5 percent of the baseline mean). Recruiters are also 8.9 percentage points more likely to recommend considering the candidate for a more senior role, relative to a baseline recommendation rate of 57.2 percent. The relatively high baseline recommendation rate reflects that recruiters were asked whether candidates should be considered for a newly available higher-level role rather than whether they would ultimately be selected for such a position.

Recruiters revise upward their estimates of candidates' current salaries (9.7 percent, $p < 0.01$) and best outside offers (9.5 percent, $p < 0.01$), consistent with counteroffers conveying information about market value. Recruiters also increase some productivity assessments: technical skills rise 2.5 percent ($p = 0.072$) and initiative rises 3.7 percent ($p < 0.01$), which are smaller in magnitude than the shifts in perceived market value.

At the same time, recruiters view these candidates as more difficult to hire. They are less likely to believe that candidates will accept the recommended offer (-0.171; $p < 0.01$), more likely to advise against accepting the counteroffer (0.31; $p < 0.01$), and more likely to reject the candidacy (0.046; $p < 0.01$). The rise in rejection is consistent with high counteroffers compressing expected surplus: perceived outside options rise substantially, increasing the wage required to close the hire, while productivity assessments increase only modestly.

4.3 Gender Differences in Returns to More Intense Negotiation

We next examine whether these belief updates differ by gender and how they map into outcomes. The model predicts that gender differences should emerge on margins that are most sensitive to the dimension along which belief updating differs across groups.

Symmetric Wage Responses. Figure 1 and Table 3 show that salary offers increase similarly for male and female candidates who negotiate intensively–approximately 10.5 percent for both groups when counteroffers exceed the target range.

These similar wage responses stem from changes in recruiters' estimates of current salaries and best outside offers. Market value estimates increase with counteroffer amounts for both groups. Male candidates experience modestly larger increases above the target range: \$2,100 in current salary estimates ($p = 0.021$) and \$1,074 in outside offer estimates ($p = 0.12$). However, these differences are economically small (1.6 percent and 0.77 percent of respective means), which suggests that outside-option updating is nearly symmetric across genders. Interestingly, we find no differential returns by gender in terms of work-from-home (WFH) recommendation, as both men and women are more likely to be offered WFH days when negotiating intensively, suggesting no differential inference about preferences along this dimension.

Divergence in Advancement and Rejection Decisions. In contrast to pay, advancement and continuation decisions diverge substantially by gender. Female candidates who negotiate intensively are 5.8 percentage points more likely to be recommended for senior-track positions, while male counterparts are 11.6 percentage points more likely. This generates a gender gap of 5.9 percentage points ($p < 0.01$; 10.6 percent of male mean below target). Female candidates also face significantly higher rejection rates relative to men, by 2.7 percentage points (31 percent of male mean; $p = 0.047$).

Figure 2 graphically illustrates the benefits and costs of negotiating more intensively, disaggregated by gender. The top panel shows the average final offers recommended by recruiters, plotted against the counteroffer amounts suggested by candidates within \$1,000 incremental cells. The final offers increase consistently with counteroffer amounts, with similar gains observed for both men and women. The bottom panels plot the likelihood of evaluators recommending senior-track positions and the probability of application rejection on the vertical axis. The bottom panels reveal stark divergence by gender. Women receive relatively favorable senior-track recommendations at lower counteroffer levels (3.6 percentage points advantage; $p = 0.019$). However, recommendations increase more steeply with counteroffer amounts for men than for women. As a result, the female advantage disappears and reverses among candidates who negotiate above the target range, where women become 2.6 percentage points less likely to receive a positive recommendation ($p = 0.076$). Similarly, women's rejection rates increase steeply with counteroffer amounts, while men's remain flat until \$155,000.

Differential Skill Updating. We examine skill assessments to understand this divergence. While female candidates receive slightly more favorable baseline assessments in some dimensions, increases in skill assessments with negotiation are driven almost entirely by male candidates. For men, higher counteroffers generate increases of 3.4 percentage points in technical skills, 3 percentage points in team skills ($p = 0.04$), and 4.5 percentage points in initiative ($p = 0.002$). These increases significantly exceed those for women: 3.0 percentage points larger for technical skills ($p = 0.1$), 3.9 percentage points for team skills ($p = 0.03$), and 2.9 percentage points for initiative ($p = 0.093$), suggesting a sizable gender gap in the productivity-related component of belief updating.

These patterns align closely with the framework. Outside-option updating is nearly symmetric across genders ($\Delta_\omega \approx 0$), while productivity-related updating differs substantially. Because wages respond to outside-option assessments and senior-track recommendations depend more directly on perceived productivity, asymmetric productivity updating translates into divergent senior-track outcomes even when wage responses remain similar across genders. The higher rejection rate for women follows a similar logic: when productivity is revised less upward, rising outside-option assessments compress expected surplus, making rejection more likely.

While these patterns are consistent with differential productivity updating, our design does not allow us to fully distinguish gender differences in productivity updating from direct penalties for violating gender norms (Rudman, 1998; Bowles, Babcock and Lai, 2007) or other forms of evaluative bias. Nonetheless, that negotiation differentially changes recruiters' skill assessments—not merely their final decisions—is itself direct evidence of differential belief updating. The structural analysis in Section 6 builds on this evidence by evaluating the extent to which this productivity-related belief updating can account for the observed gender differences in hiring outcomes.

Aggregate Gender Gaps. The divergent returns to negotiation offset in the aggregate. Women have a relative advantage at lower levels of negotiation intensity but face increasing penalties at higher levels, while men are relatively penalized for less intensive counteroffers. As a result, unconditional gender differences in salary offers, senior-track recommendations, and rejection rates are small. Appendix Table B4 reports these differences across the full sample.

4.4 Robustness Checks and Additional Data Quality Checks

4.4.1 Alternative Specifications

We estimate several alternative specifications to assess the robustness of our findings. The first specification is the baseline model without any controls. The second specification controls for recruiter-level pre-treatment characteristics observed on the platform, including gender, race, hours worked, hourly rate, education, and recruiting experience. The third specification adds biography fixed effects to absorb persistent differences across resume templates. The fourth specification replaces these recruiter controls with recruiter fixed effects. The fifth specification includes evaluation-order fixed effects, controlling for the position in which each resume was entered into the evaluation system.[4]

Tables B5 and B6 present the results. Across specifications, the estimated effects remain stable, consistent with successful randomization and indicating that the findings are not driven by observable recruiter characteristics, resume-template composition, recruiter-specific evaluation tendencies, or evaluation-order effects.

4.4.2 Incentive Compatibility and Recruiter Attentiveness

We have made a deliberate decision not to layer our experiment with incentive compatibility mechanisms—such as performance bonuses tied to evaluation accuracy—in order to avoid introducing gaming incentives where recruiters might select the candidate they perceive as most employable rather than providing their best assessment of each candidate's fit. This raises a natural question: are recruiters exercising genuine effort, or are they paying insufficient attention to the candidate materials? If the latter, our estimates risk capturing noise rather than true evaluative responses to candidate characteristics and negotiation behavior.

To address this concern, we perform two exercises. Both exploit non-experimental variation in resume content that is orthogonal to our treatment assignments.

Leadership keywords and skill evaluations. We count the number of leadership and management-related keywords appearing in each candidate's resume (e.g., *led*, *managed*, *supervised*, *coor-*

[4]Since all recruiters within a treatment arm received the same package of candidates, we do not observe the actual order in which each recruiter chose to review individual resumes. The evaluation-order fixed effects therefore capture the system entry position rather than the actual viewing sequence.

dinated, *mentored*; see Table B7 notes for the full list) and define an indicator for whether this count exceeds the sample median. If recruiters are paying attention to resume content, candidates whose resumes contain more leadership signals should receive more favorable evaluations. Table B7 supports this prediction. Candidates with above-median leadership keywords are rated significantly higher across all three skill dimensions, with the strongest effect on team skills: a 5.9 percentage point increase in the probability of receiving a positive rating ($p < 0.01$)—consistent with the content of the keywords. The effect on technical skills (2.6 percentage points, $p < 0.01$) and initiative skills (1.3 percentage points, $p < 0.10$) is positive but smaller in magnitude.

Table B8 turns to recruiter decisions. Candidates with above-median leadership keywords are 1.4 percentage points less likely to have their applications rejected ($p < 0.05$) and are more likely to receive senior-role recommendations. The latter relationship is particularly pronounced among male candidates, although the gender difference is estimated imprecisely ($p = 0.12$). We find no meaningful effects on work-from-home recommendations or salary offers.

Resume quality index and recruiter decisions. We construct a composite resume quality index based on observable candidate characteristics. Specifically, we compute the mean of three indicators: (i) whether the candidate holds a computer science degree, (ii) whether the candidate has senior-level work experience (identified by job titles containing II" or Sr."), and (iii) whether the candidate has above-median total work experience (in months). We then classify candidates as high-quality if their index exceeds the sample median. The idea is that if recruiters are engaging meaningfully with the candidate packages, their decisions should differ systematically between high- and low-quality resumes—and they do. Table B9 shows that high-quality candidates are rated significantly higher on technical skills (3.4 percentage points, $p < 0.01$) and initiative skills (2.7 percentage points, $p < 0.01$), with no significant difference on team skills. Table B10 turns to recruiter decisions. High-quality candidates receive higher salary offers ($1,281$, $p < 0.01$), are 2.9 percentage points more likely to be recommended for senior-track positions ($p < 0.01$), 2.1 percentage points more likely to receive a work-from-home recommendation ($p < 0.05$), and 2.2 percentage points less likely to have their applications rejected ($p < 0.01$).

Taken together, these two exercises provide reassuring evidence that recruiters were attentive to the candidate materials.

4.5 Heterogeneity by Candidates' Characteristics

Recruiters may differentially evaluate negotiating candidates based on other observable characteristics in their resumes. Although these characteristics are balanced across genders in our sample, returns to negotiation could still vary with attributes such as experience. Alternatively, our results could be driven by a specific subgroup of female candidates. For example, women of childbearing age may face greater penalties in senior-track recommendations.

We examine these possibilities in Table B11 by exploring heterogeneity across candidate characteristics. To proxy for age and potential fertility-related concerns, we use the graduation year reported on the CV (Panel A). We also exploit variation in years of work experience (Panel B). Overall, we find no evidence of heterogeneous effects along these dimensions, suggesting that the gender penalty associated with aggressive counteroffers is unlikely to be driven by concerns related to fertility or experience.

5 The Role of Contextual Factors in Negotiation

5.1 Heterogeneity by Recruiters

We examine how prevalent these disparities are across recruiters, using a male-favoring propensity measure, as described in Section 3.1. We find that recruiters who favor male candidates evaluate female candidates less favorably than other recruiters do. As shown in Table B12, male-favoring recruiters recommend salaries that are \$873 lower for female candidates than for male candidates when the counteroffer is below the target range (p-value = 0.11). When female candidates counteroffer above the target, the penalty increases to \$2,154 (p-value = 0.04). Recommendations for senior roles display a similar pattern, with larger penalties in magnitude among male-favoring recruiters, although the estimates are imprecise. Notably, the penalty female candidates face for counteroffering above the target in terms of outright rejection is concentrated among male-favoring recruiters.

We next examine whether the main patterns vary with the level of gender bias in the

recruiter's location. Following Kline, Rose and Walters (2022*b*), we construct MSA-level measures of gender bias using Gender-Science Implicit Association Test (IAT) scores from Harvard's Project Implicit, averaged over 2015-2020. Higher values indicate stronger implicit associations linking men with careers (rather than family) or science (rather than humanities), and women with family or humanities, respectively. Although both measures capture gender stereotyping, they may reflect different dimensions of bias relevant to labor-market evaluations. We assign each recruiter the average IAT score of their MSA and classify recruiters as residing in a high-bias MSA if their local score exceeds the sample median. The two measures are moderately correlated (correlation = 0.40).

As shown in Table B13, the differential effects of negotiation on advancement decisions are concentrated in high-bias labor markets as measured by the Gender-Career IAT. In these locations, women receive more favorable senior-role recommendations than men when counteroffers remain below the target range, but this advantage reverses when candidates negotiate above the target range ($p = 0.024$). Women who negotiate aggressively also receive lower salary recommendations (1.2 percent; $p = 0.10$). Rejection patterns are qualitatively similar, although the estimates are less precise ($p = 0.15$).

Results using the Gender-Science IAT are broadly similar. The estimated effects generally point in the same direction, though they are less precisely estimated. We also find suggestive evidence that recruiters in high-bias MSAs are more likely to recommend work-from-home flexibility for women making moderate counteroffers, but that advantage disappears and reverses at higher levels of negotiation intensity.

In terms of recruiters' own characteristics, we show that female recruiters make broadly similar assessments to male recruiters, though they recommend $1,612 less in salary offers to candidates who counteroffer beyond the target range (Table B14). The gender differences are not significantly moderated by recruiter gender.

5.2 The Spillover Effect of Other Negotiating Candidates

We also assess how evaluation patterns shift when a larger share of candidates negotiate intensively. A higher prevalence of negotiation may shift relative comparisons (Radbruch and Schiprowski, 2025; Kessler, Low and Shan, 2026), or the rationing of senior-track positions, thereby affecting how individual candidates are evaluated. Recruiters were randomly assigned

to batches in which 25 percent, 50 percent, or 75 percent of applicants submitted counteroffers exceeding the target salary range. Table B15 presents the differential effects in batches with high prevalence of intense negotiation (i.e., 50 and 75 percent batches). Overall, we find limited evidence that the share of over-target counteroffers in the batch significantly affects gender gaps in returns to negotiations. Similarly, we find no robust evidence that the share of women making aggressive counteroffers systematically alters employer responses (Table B16).

5.3 The Impact of Pay Transparency

Next, we examine the impact of pay transparency on gender differences in candidates' salary outcomes and assessments. Pay transparency could mitigate these disparities by narrowing the scope for negotiation and creating clearer expectations. However, it might also exacerbate the gap if recruiters penalize candidates, particularly women, who negotiate beyond the boundaries set by the disclosed salary range.

We introduced pay transparency to the treatment group by informing recruiters that a target salary range was included in the job advertisement. In contrast, the control group was only informed that pay would be commensurate with experience and market levels. To assess whether the pay transparency treatment was salient and whether recruiters paid attention to the message, we examine their responses in Table 4. Specifically, we test whether recruiters in the treatment group were more likely to express concerns commonly associated with pay transparency. Recruiters in the pay transparency treatment were 6.6 percentage points more likely to report concerns about potential legal consequences for making exceptions to accommodate higher salary requests from candidates exceeding the target pay range (p-value = 0.095). This aligns with the message presented to them prior to accepting the task. Additionally, treated recruiters were 6.3 percentage points more likely to mention concerns about renegotiation pressures and the need to offer higher salaries to future candidates (p-value = 0.14). These findings provide first-stage evidence that recruiters paid attention to the treatment and that the message influenced their attitudes.

We present the results for the treatment effects in Table 5. Overall, we find that the transparency treatment has minimal impact on candidates' salary offers. However, in the second column, we observe that pay transparency moderates the gender differences previously observed in senior-track recommendations. Specifically, differences in how recruiters rewarded

candidates based on gender and whether they counter-offered above the target are mitigated under transparency.

For example, among male candidates, without pay transparency, counteroffers above the target increase the likelihood of receiving a senior-track recommendation by 15.7 percentage points (p-value < 0.001). Under pay transparency, this effect is substantially attenuated. Female CVs receive 6.2 percentage points higher recommendations than male CVs when counteroffers are below the target in the absence of transparency (p-value $= 0.001$), but this differential is reduced by 5.1 percentage points under pay transparency (p-value $= 0.1$). Similarly, without pay transparency, women face an 11 percentage point penalty for negotiating beyond the target compared to men (p-value < 0.001); this gap shrinks by 9.1 percentage points under transparency (p-value $= 0.04$). We find little evidence that transparency meaningfully changes rejection decisions.

It is important to note that this does not imply the full consequence of pay transparency on salaries. Pay transparency can influence outcomes at the initial offer stage, which we abstract from in this analysis. When we presented recruiters with scenarios involving a wage counteroffer, the perceived informational advantage of candidates under pay transparency was mitigated. Recruiters could assume that candidates, even in the non-transparency treatment, had some indirect information about the wage range to inform their counteroffers. While this limitation affects our ability to fully measure the impact of pay transparency policies, our results suggest that such policies can play a role in mitigating the discretion that contributes to disparities in negotiation outcomes.

6 Structural Estimation

Section 4 shows that recruiters revise perceived skill upward for men who renegotiate aggressively but not for women, and that this asymmetry coincides with gaps in seniority recommendations and rejection. The reduced-form design cannot tell us how much of these downstream gaps the skill-belief asymmetry alone generates, since we cannot experimentally move recruiters' beliefs independently of gender and the counteroffer. We therefore estimate a structural model of the firm's hiring decision adapted from Section 1, which lets us hold beliefs fixed and ask what remains, and use the same model to pursue two further questions

the reduced form cannot reach: what counteroffer would be optimal for a candidate to make, and what pay transparency changes about the firm's decision rule.

Estimation. A candidate in cell c is summarized by recruiter beliefs $(\hat{v}_c, \hat{\omega}_c)$ — perceived skill and outside option. The firm chooses a wage w and seniority recommendation s to maximize expected profit. Beyond the parameters introduced in Section 1, we add a small number of additional terms, most of which are tuning parameters detailed in Appendix C. The one we highlight here is a further, gender- and ask-specific shifter in the firm's objective: for each gender, one constant governs how steeply the firm's willingness to grant seniority rises with counteroffers above the posted-range threshold, and another governs the response to counteroffers below it—four rates for the seniority margin, with an analogous set of four estimated jointly for wages (eight in total, denoted γ). For a female candidate counteroffering (a) above the threshold, for instance, the seniority shifter takes the form $\gamma_{fo} \times (a - \$155\text{K})$, with the analogous linear form for the other three gender-by-direction combinations. We include γ so that the belief-equalization counterfactual below can isolate what survives as the gender difference-in-differences once belief differences are turned off. We estimate the full model jointly on six survey outcomes across ten gender $\times$ counteroffer-quintile cells, matching 60 moments with 38 parameters and 22 overidentifying restrictions, and it reproduces every targeted moment within sampling error, including the gender gaps in seniority and rejection (Appendix Table C2).

Table 6 reports the selected set of estimates. Production is complementary with seniority with δ estimated to be 0.27. The two rigidities are both large and precisely estimated (κ_w around 11, κ_s around 88), and their quadratic form gives each a direct dollar interpretation: a \$1K deviation of the wage from its anchor costs about \$5.4K in lost surplus, and a \$5K deviation costs about \$135.5K. On the seniority margin, moving the recommendation by 0.1 costs about \$441, and moving across its entire range—from not recommending a candidate to strongly recommending them, $s = 0$ to $s = 1$—costs about \$44K. Because a full-range move on the seniority margin costs less than a below-typical wage move of just \$5K, wage-setting is the far more rigid margin.[5] Offered wages track the candidate's own perceived market value closely (α is around 0.91), leaving little room for the institutional anchor to pull offers away

[5]Even though the raw coefficient κ_s exceeds κ_w, the two coefficients are not on the same scale to begin with, since wages vary continuously in dollars while s is bounded on $[0, 1]$.

from perceived outside options. Candidates place a comparatively modest dollar value on the senior-track recommendation ($\tau \approx \$3$K, about 2 percent of the average wage offer). Because acceptance depends on the additively separable index $w + \tau s - \hat{\omega}$, τ can be interpreted as the wage-equivalent value of a one-unit increase in s, holding w fixed—the rate at which the candidate would trade this recommendation for salary. The modest size is consistent with what the outcome actually measures: recruiters report whether they would recommend the candidate for *consideration* for a more senior role that *may* become available, not a guaranteed promotion or placement.

Isolating the productivity belief channel. To isolate how much of each gap the belief asymmetry alone generates, we equalize beliefs across genders — replacing each counteroffer quintile's $(\hat{v}, \hat{\omega})$ with its gender-pooled average, so that a man and a woman in the same quintile face identical perceived skill and outside option — and recompute each gender difference-in-differences from the firm's first-order conditions (see Table C4). With beliefs held gender-symmetric, little of the original gap remains. A complementary exercise confirms that holding beliefs at their observed, gender-asymmetric values and instead zeroing the remaining margin of the firm's response leaves each difference-in-differences essentially at its full, observed magnitude. The gap is therefore in large part attributable to the belief asymmetry itself.

Optimal counteroffer. The estimates from Section 4 show how recruiters respond to negotiation intensity along several margins—salary offers, senior-track recommendations, and rejection. We now ask what these responses imply for a candidate's optimal counteroffer. Candidates trade off a higher hiring probability at a lower counteroffer against better compensation at a higher one. Formally, a candidate of gender g chooses the counteroffer c to maximize expected utility

$$\mathrm{EU}(c) = P_{\text{hire}}(c, g)\, u\Big(w^\star(c, g) + \tau s^\star(c, g)\Big) \tag{11}$$

where P_{hire} is the hiring probability, computed from the probability that the recruiter retains the candidate's candidacy and the candidate's implied acceptance probability at counteroffer c, and $w^\star + \tau s^\star$ is the resulting compensation if hired. We compute $c^\star = \arg\max_{c \in [\$120, \$200]\text{K}} \mathrm{EU}(c)$ under three preference specifications. The first is risk neutrality,

$u(x) = x$. The second is constant relative risk aversion, $u(x) = (x^{1-\rho} - 1)/(1-\rho)$, evaluated at $\rho = 0.5$. The third is the reference-dependent value function of Tversky and Kahneman (1992),

$$v(x) = \begin{cases} (x-r)^{\alpha}, & x \geq r \\ -\lambda\,(r-x)^{\alpha}, & x < r \end{cases}$$

with reference point $r = \$130\text{K}$ (the initial bid) and their estimated parameters $\alpha = 0.88$ and $\lambda = 2.25$ (see Appendix C.4 for details).

Figure 3 plots the resulting optimal counteroffers by gender.[6] Across all parameter combinations we consider, the counteroffer that maximizes expected utility is lower for women than for men. For women, the optimal counteroffer lies 11-12 percent above the initial offer, compared with 20-30 percent for men. This gap traces to hiring probability rather than compensation: as the counteroffer rises, a man's perceived skill, and so his surplus to the firm, erodes far less than a woman's, so men can renegotiate substantially more before hiring probability turns against them. A female candidate who correctly anticipates employer responses would therefore optimally negotiate less aggressively than an otherwise identical male candidate, resulting in a gender gap in the optimal counteroffer of between 9 and 17 percentage points across utility specifications. Table C7 further shows that a counterfactual exercise in which we equalize beliefs about productivity across genders closes these gender gaps in optimal counteroffers.

Mechanisms behind the effect of pay transparency. Finally, we re-estimate the model separately under pay secrecy and transparency to examine what drives changes in the firm's decision-making across the two regimes (Table C8). Two patterns emerge. First, the overall resistance to senior-track recommendations rises under transparency (κ_s increases, $p = 0.1$), consistent with the firm becoming more broadly cautious about senior-track recommendations for men and women alike. Second, and more directly tied to the gender gap itself, in Table C9, the belief channel barely moves across regimes—the belief-driven component of the seniority difference-in-differences is 0.037 under secrecy and 0.043 under transparency—while the gender-and-ask-specific residual shifter γ does the opposite: it accounts for the majority of the gap under secrecy (0.061, 61 percent of the total), but under transparency the belief

[6] See Table C5 for corresponding estimates.

channel alone is enough to account for essentially the entire, now much smaller and statistically insignificant total gap (0.035). Because the belief channel is essentially unchanged while γ's contribution disappears, the closing of the seniority gender gap under transparency reflects a shrinking of the firm's gender-differentiated discretion specifically, not a change in how recruiters' beliefs about skill translate into recommendations—consistent with transparency reducing disparate treatment that is not grounded in productivity beliefs.

7 Conclusion

We show that the negotiation stage contributes to gender differences in employer responses beyond the initial offer. Using a field experiment with professional recruiters, we find that negotiation behavior changes employers' beliefs about outside options and productivity, and that these beliefs map differently into pay, career-track placement, and rejection decisions. While recruiters generally evaluate female candidates more favorably when counteroffers are moderate, these advantages reverse at higher levels of negotiation intensity. When candidates negotiate above the target range, male candidates receive more favorable productivity-related updates and senior-track recommendations, while female candidates face higher rejection rates and weaker advancement recommendations. In contrast, wage responses are similar across genders, consistent with similar updating of outside-option beliefs.

More broadly, our findings highlight the value of observing both employer beliefs and employer actions. Doing so allows us to move beyond documenting differences in outcomes and shed light on the mechanisms through which negotiation affects labor-market decisions. In addition, the largest gender differences emerge not in pay-setting itself but in advancement recommendations and decisions about whether to proceed with a candidacy. Focusing solely on wage outcomes may therefore understate the role of negotiation in generating inequality and highlights the importance of examining multiple margins of employer decision-making.

An important implication of our findings concerns not just how negotiation shapes immediate hiring outcomes, but how it can distort the evaluation of skill itself. Because skill assessments formed at the point of hire are likely to carry forward into subsequent promotion and talent-management decisions, such biased assessments have the potential to compound over time—amplifying inefficient talent allocation within firms and generating diverging career

trajectories for candidates of comparable underlying ability (Benson, Li and Shue, 2019).

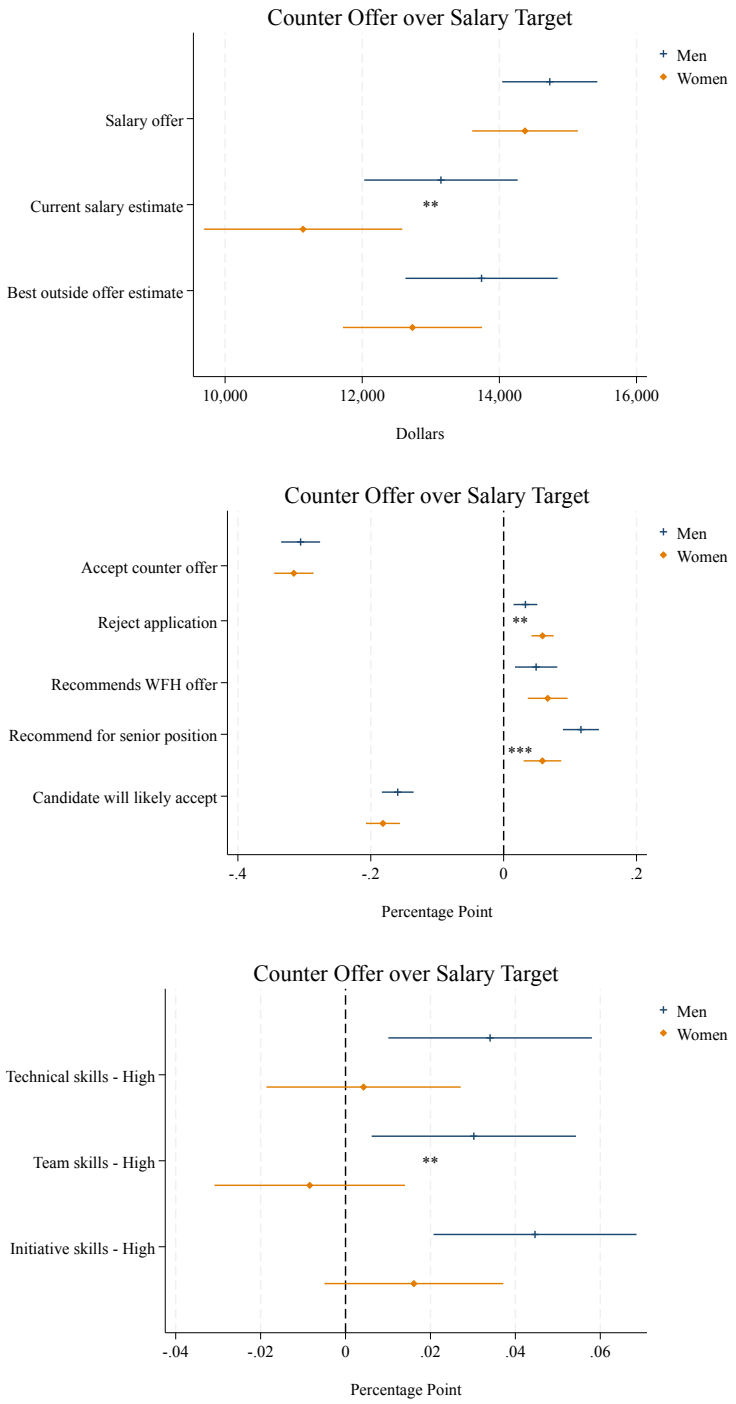
Our findings also carry important policy implications. First, because the largest gender differences emerge in advancement and selection decisions rather than wage-setting itself, efforts to reduce disparities may need to focus not only on salary-setting processes but also on the broader evaluation systems that shape promotion and hiring decisions.

Second, the returns to negotiation depend not only on workers' behavior but also on how employers interpret that behavior. Interventions that focus exclusively on encouraging negotiation or providing negotiation training may therefore be insufficient if employer responses remain unchanged.

Third, we find that pay transparency attenuates gender differences in advancement decisions. Although these outcomes are not directly targeted by transparency policies, they appear to be affected through broader constraints on employer discretion. This suggests that institutional features of the wage-setting environment can influence not only pay but also other dimensions of employer decision-making. Future research should further examine how transparency, evaluation systems, and other forms of institutional design shape negotiation processes and their implications for inequality.

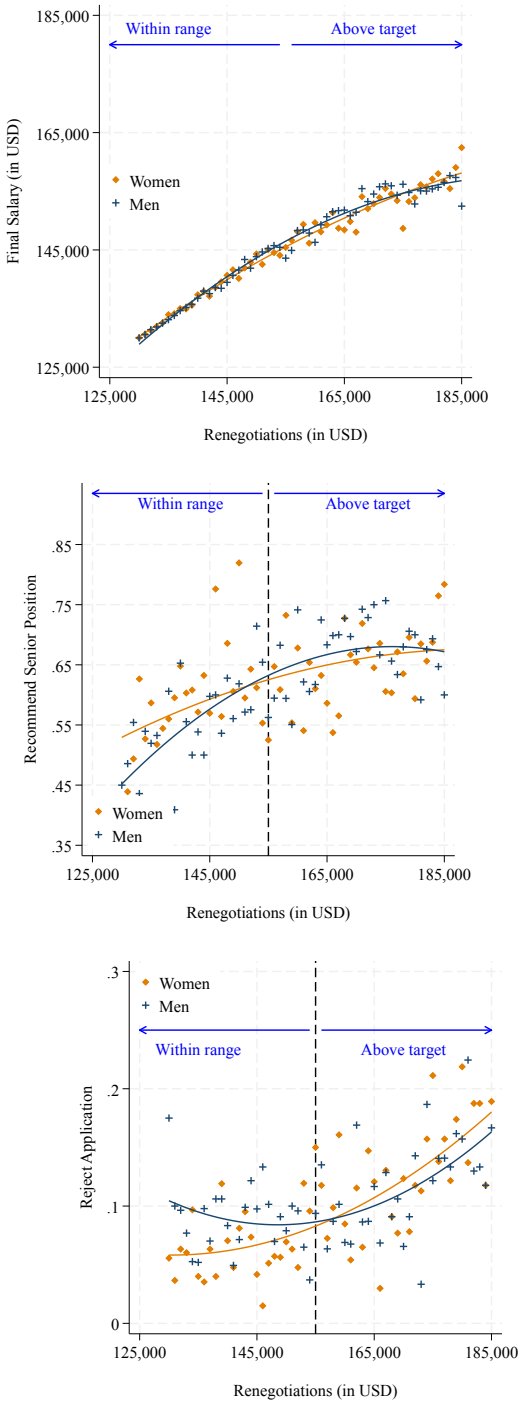
Figures and Tables

FIGURE 1. The Effects of Intensity of Negotiation by Gender



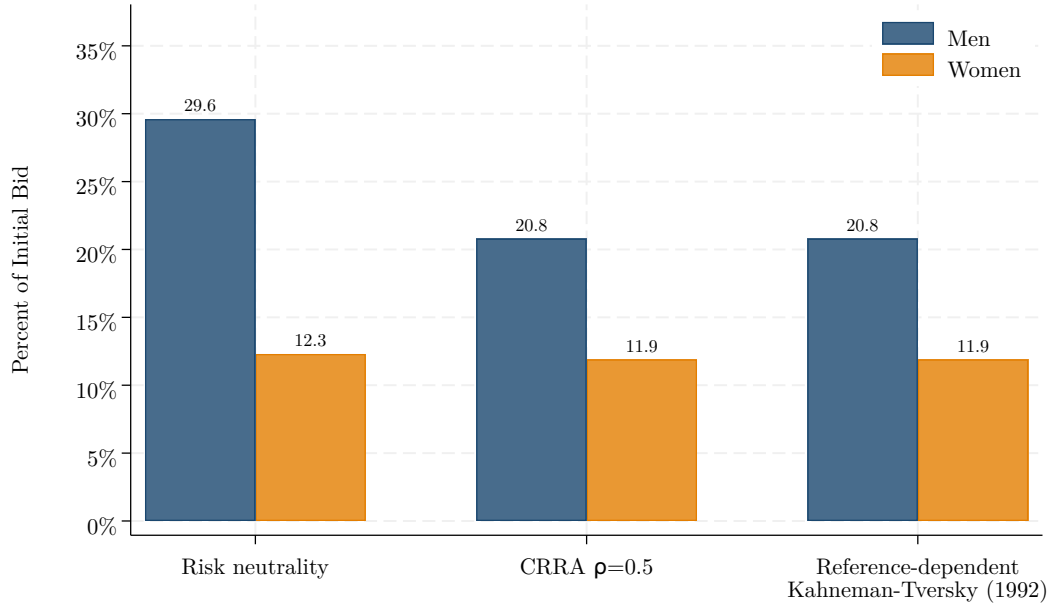
Notes: This figure displays the coefficients for the indicator variable representing whether candidates' counteroffers exceeded the target salary range, separately by gender. We also report the statistical significance of the interaction between this variable and the female indicator, showing significance levels as follows: * p < 10%, ** p < 5%, and *** p < 1%.

FIGURE 2. Benefits and Costs of Intense Negotiation



Notes: This figure illustrates the relationship between hypothetical counteroffers made by fictitious candidates and various perceptions and decisions by recruiters, separately by gender. The figure is organized clockwise from the top left: final salary offer recommendations, whether recruiters made senior-track recommendations, and whether the candidate's application was rejected. To create the graphs, we calculated the mean of the outcome variable within each $1,000 increment of the counteroffer range.

FIGURE 3. Optimal Counteroffer by Gender



Notes: This figure plots the optimal counteroffer $c^\star$ in percent of the \$130K initial bid maximizing the candidate's expected utility against the estimated firm from Equation (11), by gender and across utility specifications: risk-neutrality, CRRA ρ=0.5, and a reference-dependent value function of Tversky and Kahneman (1992) with reference point r=\$130K and estimated parameters $\alpha = 0.88$ and $\lambda = 2.25$. See Table C5 for corresponding table.

Table 1 – Summary Statistics - Recruiters

Variables	Mean	SD	Min	Max	N
Recruiter Profiles					
Female	0.736	0.441	0	1	493
White	0.533	0.499	0	1	493
Black	0.247	0.432	0	1	493
Total Hours	366.207	1,266.645	0	11,294	493
Hourly Rate	64.130	30.709	10	150	493
Bachelor degree	0.892	0.310	0	1	493
5 years of experience or more	0.769	0.422	0	1	493
Resume characteristics					
Female CV	0.505	0.057	0	1	493
Total Experince (Month)	53.761	5.334	19	70.4	493
Counteroffer over target	0.514	0.211	.143	1	493
Percent over target: 25	0.325	0.469	0	1	493
Percent over target: 50	0.310	0.463	0	1	493
Percent over target: 75	0.365	0.482	0	1	493
Pay Transparency Advert	0.513	0.500	0	1	493
Recruiter Opinions on Counteroffers					
Counteroffers over target are common	0.911	0.285	0	1	493
Counteroffers over target for women are common	0.615	0.487	0	1	493
Counteroffer is informative about current salary	0.621	0.486	0	1	493
Counteroffer is informative about outside options	0.548	0.498	0	1	493
Counteroffer is informative about personality	0.430	0.496	0	1	493
Recruiter Concerns on Accommodating Above-Range Salary Requests					
Wider re-negotiations; higher salaries to future candidates	0.663	0.473	0	1	493
Considerations of fairness to other candidates	0.647	0.478	0	1	493
Concerns about legal consequences for making exceptions	0.264	0.441	0	1	493

Notes: This table presents summary statistics for the 493 recruiters hired for this study. The first panel provides information on recruiter demographics and job experience, which are used as control variables in subsequent analyses. The second panel presents recruiter-level averages based on the 16 resumes each recruiter evaluated. This includes the percentage of female candidates, the total work experience listed on resumes, the proportion of resumes with counteroffers above the salary range, and whether the packet included 25, 50, or 75 percent of candidates submitting counteroffers more than the target range. The table also includes summary statistics indicating whether the recruiter's packet was assigned to the pay transparency regime. The third panel summarizes responses from post-experimental debriefing questions. These include whether counteroffers above the target salary range are common overall and specifically for female candidates, as well as how informative counteroffers are perceived to be. Responses were coded into a dummy variable indicating answers above the midpoint on a 4-point scale. Additionally, recruiters were asked about concerns related to accommodating above-range salary requests, with dummy variables coded for whether each concern was selected.

Table 2 – Effect of Counteroffer Over Salary Range

Variables	Below Target (1)	Over Target (2)	Difference (2)-(1)	p-value
Recommendations				
Salary offer	138,184	152,775	14,567.15*** (390.60)	0.000
Accept counter offer	0.556	0.247	-0.309*** (0.015)	0.000
Reject application	0.077	0.121	0.046*** (0.008)	0.000
Recommends WFH offer	0.586	0.644	0.058*** (0.016)	0.000
Recommend for senior position	0.572	0.661	0.089*** (0.013)	0.000
Market Value and Options				
Current salary estimate	125,237	137,407	12,152.54*** (654.74)	0.000
Best outside offer estimate	139,606	152,907	13,249.12*** (553.14)	0.000
Candidate will likely accept	0.893	0.723	-0.171*** (0.012)	0.000
Skill Assessments				
Technical skills - High	0.800	0.821	0.020* (0.011)	0.072
Team skills - High	0.806	0.818	0.011 (0.011)	0.321
Initiative skills - High	0.814	0.844	0.030*** (0.011)	0.004
N			7,749	
Number of clusters			493	

Notes: This table examines how recruiters perceive and evaluate candidates who request counteroffers above the target salary range compared to those who request below it. The first section presents recommendations on the final salary offer, whether to accept the candidate's counteroffer, reject the application altogether, offer work-from-home arrangements, and recommend the candidate for a new senior-track position. The second section reports recruiters' estimates of the candidate's market value and likelihood of accepting the offer. The third section provides results for recruiters' assessments of candidates' skills. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table 3 – Effect of Counteroffer Over Salary Range By Gender

	Female CV			Male CV			Difference F-M	
	Below Target	Difference	*p*-value	Below Target	Difference	*p*-value	Difference	*p*-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Recommendations								
Salary offer	138,123	14,374.08***	0.000	138,247	14,734.91***	0.000	-379	0.376
		(467.43)			(420.78)		(428.08)	
Accept counter offer	0.568	-0.316***	0.000	0.544	-0.306***	0.000	-0.009	0.659
		(0.018)			(0.018)		(0.021)	
Reject application	0.066	0.059***	0.000	0.087	0.033***	0.003	0.027**	0.047
		(0.010)			(0.011)		(0.013)	
Recommends WFH offer	0.599	0.066***	0.000	0.572	0.049**	0.012	0.017	0.413
		(0.018)			(0.019)		(0.021)	
Recommend for senior position	0.590	0.058***	0.001	0.554	0.116***	0.000	-0.059***	0.005
		(0.017)			(0.016)		(0.021)	
Market Value and Options								
Current salary estimate	125,513	11,137.54***	0.000	124,951	13,147.81***	0.000	-2,100**	0.021
		(878.06)			(679.04)		(905.40)	
Best outside offer estimate	139,677	12,732.72***	0.000	139,534	13,740.06***	0.000	-1,074	0.117
		(615.30)			(673.44)		(683.45)	
Candidate will likely accept	0.897	-0.182***	0.000	0.890	-0.159***	0.000	-0.023	0.175
		(0.015)			(0.015)		(0.017)	
Skill Assessments								
Technical skills - High	0.810	0.004	0.760	0.790	0.034**	0.020	-0.030	0.101
		(0.014)			(0.015)		(0.018)	
Team skills - High	0.823	-0.008	0.536	0.788	0.030**	0.039	-0.039**	0.033
		(0.014)			(0.015)		(0.018)	
Initiative skills - High	0.832	0.016	0.208	0.795	0.045***	0.002	-0.029*	0.093
		(0.013)			(0.015)		(0.017)	
N		3,905			3,844			

Notes: This table examines how recruiters perceive and evaluate candidates who request counteroffers above the target salary range compared to those who request below it, with differential effects by gender. The first and second panels provide separate results for male and female candidates, while the third panel tests the difference in outcomes between genders. The first section presents recommendations on the final salary offer, whether to accept the candidate's counteroffer, reject the application altogether, offer work-from-home arrangements, and recommend the candidate for a new senior-track position. The second section reports recruiters' estimates of the candidate's market value and likelihood of accepting the offer. The third section provides results for recruiters' assessments of candidates' skills. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table 4 – Effect of Pay Transparency on Recruiter Attitudes

Variables	Effect of Pay Transparency	p-value
Informational value of negotiations		
Quality of the pool of candidates: High	0.008	0.847
	(0.043)	
Counteroffer is informative about current salary	0.045	0.311
	(0.044)	
Counteroffer is informative about outside options	0.030	0.505
	(0.045)	
Counteroffer is informative about personality	0.048	0.288
	(0.045)	
Prevalence of negotiations		
Counteroffers over target are common	-0.021	0.410
	(0.026)	
Counteroffers over target for women are common	0.000	0.993
	(0.044)	
Concerns regarding pay transparency		
Fear of wider re-negotiations & need of higher salaries to future candidates	0.063	0.142
	(0.043)	
Considerations of fairness to other candidates	0.045	0.295
	(0.043)	
Concerns about potential legal consequences for making exceptions	0.066*	0.095
	(0.040)	
N	493	

Notes: This table shows results from the post-experimental survey where we asked recruiters about various aspects of negotiation. In the pay transparency treatment, recruiters were informed that the job advertisement disclosed the salary range. In contrast, the control group was told that the advertisement only stated the salary would be commensurate with experience. The variables include whether counteroffers above the target salary range are common overall and specifically for female candidates, as well as how informative counteroffers are perceived to be. Responses were coded into a dummy variable indicating answers above the midpoint on a 4-point scale. Additionally, recruiters were asked about concerns related to accommodating above-range salary requests, with dummy variables coded for whether each concern was selected. Robust standard errors are used.

Table 5 – Effects of Pay Transparency

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
Pay transparency	584.119	0.018	0.015	-0.005
	(420.194)	(0.032)	(0.017)	(0.038)
Over target	15076.394***	0.157***	0.034**	0.058**
	(624.080)	(0.023)	(0.015)	(0.027)
Over target × Pay transparency	-529.073	-0.070**	-0.008	-0.017
	(848.435)	(0.034)	(0.022)	(0.038)
Female CV	118.037	0.062***	-0.023**	0.028
	(301.013)	(0.021)	(0.011)	(0.019)
Female CV × Pay transparency	-473.673	-0.051*	0.005	-0.003
	(436.134)	(0.031)	(0.017)	(0.027)
Over target × Female CV	-877.686	-0.110***	0.026	0.029
	(622.958)	(0.031)	(0.018)	(0.031)
Over target × Female CV × Pay transparency	797.866	0.091**	0.004	-0.022
	(885.844)	(0.044)	(0.027)	(0.042)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749

Notes: This table examines how recruiters evaluate candidates under pay transparency treatment compared to the control group, with differential effects by gender. In the pay transparency treatment, recruiters were informed that the job advertisement disclosed the salary range. In contrast, the control group was told that the advertisement only stated the salary would be commensurate with experience. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table 6 – Structural parameter estimates

Parameter	Estimate	95% CI
Production and rigidity		
δ	0.2737	[+0.2592, +0.2864]
κ_w	10.84	[+9.43, +11.93]
κ_s	88.19	[+85.50, +88.44]
α	0.91	[+0.89, +0.92]
τ	2.98	[+2.77, +3.37]
c_R	4.01	[+3.17, +4.80]
N		7,733

Notes: The table reports selected parameters of the structural model described in Section 6; details of the estimation procedure are given in Appendix C. The estimation sample is the full experimental sample, $N = 7{,}733$. 95% confidence intervals are obtained from a recruiter-clustered bootstrap ($B = 200$). We report $Q(\hat{\theta}) = 107.10$, the value of the minimum-distance criterion at the estimated parameters—the inverse-variance-weighted sum of squared gaps between the model-implied and sample moments. A cluster-robust reweighting gives $Q = 68.47$, and a sensitivity analysis shows the residual misfit is orthogonal to the paper's decomposition (Table C3).

References

Adler, Laura. 2020. "What's a Job Candidate Worth? Explaining Variation in Pay-Setting Practices." Working Paper.

Agan, Amanda Y, Bo Cowgill, and Laura K Gee. 2021. "Salary history and employer demand: Evidence from a two-sided audit." National Bureau of Economic Research.

Alan, Manning. 2011. "Imperfect competition in the labor market." In *Handbook of labor economics*. Vol. 4, 973–1041. Elsevier.

Amanatullah, Emily T, and Michael W Morris. 2010. "Negotiating gender roles: Gender differences in assertive negotiating are mediated by women?s fear of backlash and attenuated when negotiating on behalf of others." *Journal of personality and social psychology*, 98(2): 256.

Amer, Abdelrahman, Ashley Craig, and Clémentine Van Effenterre. 2024. "Decoding Gender Bias: The Role of Personal Interaction."

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *The Quarterly Journal of Economics*, 132(4): 1553–1592.

Arrow, Kenneth, J. 1973. "The Theory of Discrimination." *Discrimination in Labor Markets*, 3.

Babcock, Linda, Michele Gelfand, Deborah Small, and Heidi Stayn. 2013. "Gender differences in the propensity to initiate negotiations." In *Social psychology and economics*. 239–259. Psychology Press.

Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. 2023. "Pay transparency and the gender gap." *American Economic Journal: Applied Economics*, 15(2): 157–183.

Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. 2022. "Do firms respond to gender pay gap transparency?" *The Journal of Finance*, 77(4): 2051–2091.

Benson, Alan, Danielle Li, and Kelly Shue. 2019. "Promotions and the peter principle." *The Quarterly Journal of Economics*, 134(4): 2085–2134.

Benson, Alan, Danielle Li, and Kelly Shue. 2024. "Potential and the Gender Promotions Gap." *SSRN Electronic Journal*.

Bertrand, Marianne, and Esther Duflo. 2017. "Field experiments on discrimination." In *Handbook of Economic Field Experiments*. Vol. 1, 309–393. Elsevier.

Biasi, Barbara, and Heather Sarsons. 2022. "Flexible wages, bargaining, and the gender gap." *The Quarterly Journal of Economics*, 137(1): 215–266.

Black, Ines, Sharique Hasan, and Rembrand Koning. 2020. "Hunting for talent: Firm-driven labor market search in America." SSRN Electronic Journal.

Bowles, Hannah Riley, Linda Babcock, and Lei Lai. 2007. "Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask." *Organizational Behavior and Human Decision Processes*, 103(1): 84–103.

Brenzel, Hanna, Hermann Gartner, and Claus Schnabel. 2014. "Wage bargaining or wage posting? Evidence from the employers' side." *Labour Economics*, 29: 41–48.

Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 39(2): 257–273.

Burn, Ian, and Kyle Kettler. 2019. "The more you know, the better you?re paid? Evidence from pay secrecy bans for managers." *Labour Economics*, 59: 92–109.

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

Caldwell, Sydnee, and Oren Danieli. 2024. "Outside Options in the Labour Market." *Review of Economic Studies*.

Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining. 2023. "Bargaining and Inequality in the Labor Market." Tech. rep., Working Paper.

Cappelli, Peter. 2019. "Your approach to hiring is all wrong." *Harvard Business Review*, 97(3): 48–58.

Card, David. 2022. "Who set your wage?" *American Economic Review*, 112(4): 1075–1090.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women." *The Quarterly journal of economics*, 131(2): 633–686.

Chotiputsilp, Brighton, and Taeho Kim. 2023. "The Cost of Salary Negotiation : An Experimental Evaluation with Job Seekers in Thailand." Working Paper.

Cornerstone OnDemand. 2024. "Manage Candidates: Add/Move to Requisition." Accessed: 2026-07-21.

Cowgill, Bo, Jonathan Davis, B Pablo Montagnes, and Patryk Perkowski. 2021. "Matchmaking principals: Theory and evidence from internal labor markets." Columbia Business School Working Paper.

Cullen, Zoë. 2024. "Is pay transparency good?" *Journal of Economic Perspectives*, 38(1): 153–180.

Cullen, Zoë B, and Bobak Pakzad-Hurson. 2023. "Equilibrium effects of pay transparency." *Econometrica*, 91(3): 765–802.

Cullen, Zoë B., Bobak Pakzad-Hurson, and Ricardo Perez-Truglia. 2025. "Pushing the Envelope: The Effects of Salary Negotiations." *NBER Working paper series*.

Duchini, E., S. Simion, and A. Turrell. 2022. "Pay transparency and gender equality." *Working Paper No. 482.*

Exley, Christine L, Muriel Niederle, and Lise Vesterlund. 2020. "Knowing when to ask: The cost of leaning in." *Journal of Political Economy*, 128(3): 816–854.

Haegele, Ingrid. 2022. "Talent hoarding in organizations." arXiv preprint arXiv:2206.15098.

Hall, Robert E, and Alan B Krueger. 2012. "Evidence on the incidence of wage posting, wage bargaining, and on-the-job search." *American Economic Journal: Macroeconomics*, 4(4): 56–67.

Harmon, N., and S. Caldwell. 2019. "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks." Working Paper.

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li. 2018. "Discretion in hiring." *The Quarterly Journal of Economics*, 133(2): 765–800.

Jäger, Simon, Benjamin Schoefer, Samuel Young, and Josef Zweimüller. 2020. "Wages and the Value of Nonemployment." *The Quarterly Journal of Economics*, 135(4): 1905–1963.

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer. 2024. "Worker beliefs about outside options." *The Quarterly Journal of Economics*, qjae001.

Kessler, Judd B., Corinne Low, and Colin D. Sullivan. 2019. "Incentivized resume rating: Eliciting employer preferences without deception." *American Economic Review*, 109(11): 3713–3744.

Kessler, Judd B, Corinne Low, and Xiaoyue Shan. 2026. "Lowering the playing field: Discrimination through sequential spillover effects." *Review of Economics and Statistics*, 108(2): 504–513.

Kline, Patrick, and Christopher Walters. 2021. "Reasonable Doubt: Experimental Detection of Job-Level Employment Discrimination." *Econometrica*, 89(2): 765–792.

Kline, Patrick, Evan K Rose, and Christopher R Walters. 2022*a*. "Systemic Discrimination Among Large US Employers." *Quarterly Journal of Economics*, 137(4): pp 1963–2036.

Kline, Patrick, Evan K Rose, and Christopher R Walters. 2022*b*. "Systemic Discrimination Among Large U.S. Employers." *The Quarterly Journal of Economics*, 137(4): 1963–2036.

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A Woodbury. 2022. "Wage posting or wage bargaining? a test using dual jobholders." *Journal of Labor Economics*, 40(S1): S469–S493.

Landay, Karen, and Sarah DeArmond. 2018. "Recruitment Process Outsourcing and Recruiter and Hiring Firm Characteristics." *Journal of Personnel Psychology.*

Leibbrandt, Andreas, and John A List. 2015. "Do women avoid salary negotiations? Evidence from a large-scale natural field experiment." *Management Science*, 61(9): 2016–2024.

List, John A. 2020. "Non est disputandum de generalizability? A glimpse into the external validity trial." National Bureau of Economic Research.

Manning, A. 2011. "Imperfect competition in the labor market." *Handbook of Labor Economics*, 4: 973–1041.

Mas, Alexandre. 2017. "Does transparency lead to pay compression?" *Journal of Political Economy*, 125(5): 1683–1721.

Minni, Virginia, Kieu-Trang Nguyen, Heather Sarsons, and Carla Srebot. 2026. "Managers and the Cultural Transmission of Gender Norms." National Bureau of Economic Research.

Morchio, Iacopo, and Christian Moser. 2023. "The Gender Pay Gap: Micro Sources and Macro Consequences." SSRN Working Paper.

Niederle, Muriel, and Lise Vesterlund. 2007. "Do women shy away from competition? Do men compete too much?" *The quarterly journal of economics*, 122(3): 1067–1101.

Palladino, Marco G, Antoine Berthou, Alexander Hijzen, Astrid Kunze, Cesar Barreto, Dogan Gülümser, Marta Lachowska, Anne Sophie Lassen, Salvatore Lattanzio, Benjamin Lochner, et al. 2025. "Firms and the gender wage gap: A comparison of eleven countries." FRB of Chicago Working Paper.

PeopleScout. 2023. "Modular RPO for Healthcare Recruitment." Accessed: 2026-03-20.

Phelps, Edmund S. 1972. "The statistical theory of racism and sexism." *The American Economic Review*, 62(4): 659–661.

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

Radbruch, Jonas, and Amelie Schiprowski. 2025. "Interview sequences and the formation of subjective assessments." *Review of Economic Studies*, 92(2): 1226–1256.

Robert Walters. 2023. "Modular RPO." Accessed: 2026-03-20.

Roussille, N. 2020. "The Central Role of the Ask Gap in Gender Pay Inequality." *Working Paper*, 1–90.

Rudman, Laurie A. 1998. "Self-promotion as a Risk Factor for Women: The Costs and Benefits of Counterstereotypical Impression Management."

Sarsons, Heather. 2019. "Interpreting Signals in the Labor Market: Evidence from Medical Referrals."

Säve-Söderbergh, Jenny. 2019. "Gender gaps in salary negotiations: Salary requests and starting salaries in the field." *Journal of Economic Behavior and Organization*, 161: 35–51.

Seitz, S., A. Gulyas, and S. Sinha. 2022. "Does pay transparency affect the gender wage gap? Evidence from Austria." *American Economic Journal: Economic Policy*.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and uncertainty*, 5(4): 297–323.

(For Online Publication)

Appendix to

Negotiation, Employer Demand, and Gender Disparities

Ratchanon Chotiputsilp, Taeho Kim, Clémentine Van Effenterre

List of Appendices

Appendix A: Conceptual Framework **A-2**

Appendix B: Additional Results **A-15**

Appendix C: Structural Estimation **A-32**

Appendix D: Recruitment Materials **A-51**

Appendix F: Questionnaires to Recruiters **A-58**

A Conceptual Framework

A.1 Beliefs: Informational Value of Negotiation

Posterior means. Define the signal loading vector

$$a_g \equiv \begin{pmatrix} \beta_{v,g} \\ \beta_{\omega,g} \end{pmatrix}, \qquad \mu_g \equiv \begin{pmatrix} \mu_{v,g} \\ \mu_{\omega,g} \end{pmatrix}.$$

Then the signal can be written as $n = a_g' z + \varepsilon$ where $z = (v, \omega)'$. The posterior mean is

$$\mathbb{E}[z \mid n, g] = \mu_g + \Sigma_g a_g \left(a_g' \Sigma_g a_g + \sigma_\varepsilon^2 \right)^{-1} \left(n - a_g' \mu_g \right). \tag{A.1}$$

Define the total signal variance

$$V_g \equiv a_g' \Sigma_g a_g + \sigma_\varepsilon^2 = \beta_{v,g}^2 \sigma_{v,g}^2 + 2\beta_{v,g}\beta_{\omega,g}\sigma_{v\omega,g} + \beta_{\omega,g}^2 \sigma_{\omega,g}^2 + \sigma_\varepsilon^2. \tag{A.2}$$

Expanding (A.1) componentwise yields the posterior means:

$$\hat{v}_g(n) \equiv \mathbb{E}[v \mid n, g] = \mu_{v,g} + \frac{\beta_{v,g}\sigma_{v,g}^2 + \beta_{\omega,g}\sigma_{v\omega,g}}{V_g} \left(n - a_g' \mu_g \right), \tag{A.3}$$

$$\hat{\omega}_g(n) \equiv \mathbb{E}[\omega \mid n, g] = \mu_{\omega,g} + \frac{\beta_{v,g}\sigma_{v\omega,g} + \beta_{\omega,g}\sigma_{\omega,g}^2}{V_g} \left(n - a_g' \mu_g \right). \tag{A.4}$$

Slopes: How negotiation updates beliefs. The key comparative statics are the slopes of posterior means in n:

$$\frac{d\hat{v}_g}{dn} = \frac{\beta_{v,g}\sigma_{v,g}^2 + \beta_{\omega,g}\sigma_{v\omega,g}}{V_g}, \tag{A.5}$$

$$\frac{d\hat{\omega}_g}{dn} = \frac{\beta_{v,g}\sigma_{v\omega,g} + \beta_{\omega,g}\sigma_{\omega,g}^2}{V_g}. \tag{A.6}$$

Both slopes are positive under standard assumptions ($\sigma_{v,g}^2, \sigma_{\omega,g}^2 > 0$ and $\sigma_{v\omega,g}$ not too negative). The critical observation is that *both slopes can differ across groups*: negotiation may raise inferred skill more for men ($d\hat{v}_m/dn > d\hat{v}_f/dn$) while simultaneously raising inferred

outside options more for women ($d\hat{\omega}_f/dn > d\hat{\omega}_m/dn$), or vice versa.

A.2 Firm's Problem

Expected profit integrates over the posterior:

$$\Pi(w,s \mid n,g) = \iint \Big[\Lambda(w+\tau s-\omega)\, S(v,w,s) - \Big(1-\Lambda(w+\tau s-\omega)\Big)c_R\Big]\, f_g(v,\omega \mid n)\, dv\, d\omega. \quad \text{(A.7)}$$

$S(v,w,s) = v(1+\delta s) - w - C_w(w) - C_s(s)$ is the net surplus from acceptance. Letting $\mathbb{E}_{n,g}[\cdot]$ denote expectation under $f_g(\cdot \mid n)$,

$$\Pi(w,s \mid n,g) = \mathbb{E}_{n,g}\Big[\Lambda(x)\, S(v,w,s) - (1-\Lambda(x))c_R\Big] \quad \text{(A.8)}$$

Define the value of acceptance relative to rejection

$$B(v,w,s) \equiv S(v,w,s) + c_R \quad \text{(A.9)}$$

Then

$$\Pi(w,s \mid n,g) = \mathbb{E}_{n,g}\Big[\Lambda(x)\, B(v,w,s)\Big] - c_R \quad \text{(A.10)}$$

so the constant $-c_R$ drops out of the choice problem.

A.3 Optimality Conditions and Comparative Statics

We derive first-order conditions by differentiating (A.8). All derivatives below are justified by standard regularity conditions, which hold here because Λ is smooth and bounded and the posterior is Gaussian.

Denote

$$\Lambda'(x) = \Lambda(x)(1-\Lambda(x)) > 0, \qquad \Lambda''(x) = \Lambda'(x)(1-2\Lambda(x)).$$

A.3.1 Wage Choice and Responses

Wage FOC (full expectation). Differentiate Π w.r.t. w. Since $\partial x/\partial w = 1$ and

$$\frac{\partial B}{\partial w} = -1 - \kappa_w(w - w_g^{\text{eff}}),$$

we obtain

$$0 = \frac{\partial \Pi}{\partial w} = \mathbb{E}_{n,g}\Big[\Lambda'(x)\, B(v,w,s) + \Lambda(x)\,(-1 - \kappa_w(w - w_g^{\text{eff}}))\Big]. \tag{A.11}$$

Rearranging yields:

$$\mathbb{E}_{n,g}\Big[\Lambda'(x)\, B(v,w,s)\Big] = \mathbb{E}_{n,g}\Big[\Lambda(x)\Big]\,\Big[1 + \kappa_w(w - w_g^{\text{eff}})\Big] \tag{A.12}$$

The left side is the expected marginal acceptance gain from raising the wage (weighted by the value of closing the hire, B). The right side is the expected marginal cost: the direct wage bill plus the adjustment cost penalty from deviating from the effective benchmark w_g^{eff}.

Fix s and define

$$F(w; n, g, s) \equiv \mathbb{E}_{n,g}\Big[\Lambda'(x)\, B(v,w,s)\Big] - \mathbb{E}_{n,g}\Big[\Lambda(x)\Big]\,\Big[1 + \kappa_w(w - w_g^{\text{eff}})\Big]. \tag{A.13}$$

Then $w^*(n,g,s)$ solves $F(w; n, g, s) = 0$. Assume the one-dimensional wage problem is locally concave so that

$$F_w(w^*; n, g, s) = \frac{\partial^2 \Pi}{\partial w^2}(w^*, s \mid n, g) < 0, \tag{A.14}$$

which ensures a locally unique wage best response.

Impact of $\hat{\omega}$ on the wage FOC. We differentiate with respect to the posterior mean parameter $\hat{\omega}_g(n)$, holding the posterior *shape* (variance and correlation structure) fixed.

This is equivalent to writing $\omega = \hat{\omega} + \tilde{\omega}$ with $\tilde{\omega}$ mean-zero under the posterior, so a change in $\hat{\omega}$ shifts x one-for-one:

$$x = w + \tau s - (\hat{\omega} + \tilde{\omega}) \quad \Rightarrow \quad \frac{\partial x}{\partial \hat{\omega}} = -1.$$

$B(v,w,s)$ does not depend on ω, hence does not depend on $\hat{\omega}$ directly:

$$\frac{\partial B}{\partial \hat{\omega}} = 0.$$

Differentiate F in (A.13) with respect to $\hat{\omega}$ by differentiating inside the expectation:

$$\begin{aligned} F_{\hat{\omega}} &= \mathbb{E}_{n,g}\Big[\Lambda''(x)\,\frac{\partial x}{\partial \hat{\omega}}\,B(v,w,s)\Big] - \mathbb{E}_{n,g}\Big[\Lambda'(x)\,\frac{\partial x}{\partial \hat{\omega}}\Big]\,\Big[1+\kappa_w(w-w_g^{\text{eff}})\Big] + \mathbb{E}_{n,g}\Big[\Lambda(x)\Big]\,\kappa_w\,\frac{\partial w_g^{\text{eff}}}{\partial \hat{\omega}} \\ &= \mathbb{E}_{n,g}\Big[-\Lambda''(x)\,B(v,w,s)\Big] + \mathbb{E}_{n,g}\Big[\Lambda'(x)\Big]\,\Big[1+\kappa_w(w-w_g^{\text{eff}})\Big] + \alpha\,\kappa_w\,\mathbb{E}_{n,g}\Big[\Lambda(x)\Big]. \end{aligned} \tag{A.15}$$

By the implicit function theorem,

$$\frac{\partial w^*}{\partial \hat{\omega}} = -\frac{F_{\hat{\omega}}}{F_w} \qquad \text{with} \qquad F_w < 0. \tag{A.16}$$

Hence the sign is determined by $F_{\hat{\omega}}$:

$$\operatorname{sign}\left(\frac{\partial w^*}{\partial \hat{\omega}}\right) = \operatorname{sign}(F_{\hat{\omega}}).$$

Sufficient condition for positive wage response to outside options. A convenient sufficient condition for $F_{\hat{\omega}} > 0$ is that the posterior mass lies in an interior region where acceptance is reasonably likely ($\Lambda(x) \geq 1/2$) and the match generates positive surplus ($B > 0$).

Since

$$\Lambda''(x) = \Lambda'(x)\Big(1 - 2\Lambda(x)\Big),$$

we have $\Lambda''(x) \leq 0$ whenever $\Lambda(x) \geq 1/2$. In that region,

$$-\Lambda''(x)B \geq 0 \quad \text{whenever } B > 0.$$

Thus the first term in $F_{\hat{\omega}}$ is nonnegative and strictly positive when both conditions hold with positive probability.

The second term,

$$\mathbb{E}_{n,g}\Big[\Lambda'(x)\Big]\,[1+\kappa_w(w-w_g^{\text{eff}})],$$

is positive when $w > w_g^{\text{eff}}$. Moreover, the additional term $\alpha\kappa_w\,\mathbb{E}[\Lambda(x)] > 0$ holds unconditionally for $\alpha > 0$. Therefore $F_{\hat{\omega}} > 0$ and hence

$$\frac{\partial w^*}{\partial \hat{\omega}} > 0.$$

An increase in $\hat{\omega}$ creates a mechanical reduction in acceptance probability for any fixed wage: $\partial \mathbf{\Lambda}(x)/\partial \hat{\omega} < 0$. To restore acceptance and avoid rejection costs, the firm must raise w. This is a first-order effect that operates regardless of skill assessments.

Impact of $\hat{v}$ on the wage FOC. Analogously, $\hat{v}$ is the posterior mean parameter of v (holding posterior shape fixed). We write $v = \hat{v} + \tilde{v}$, so $\partial v/\partial \hat{v} = 1$. x does not depend on v, so $x_{\hat{v}} = 0$. But $B = S + c_R$ depends on v linearly through output $v(1+\delta s)$, hence

$$\frac{\partial B}{\partial \hat{v}} = \frac{\partial}{\partial \hat{v}}\Big(v(1+\delta s)\Big) = 1 + \delta s.$$

Differentiating (A.13) with respect to $\hat{v}$:

$$F_{\hat{v}} = \mathbb{E}_{n,g}\Big[\mathbf{\Lambda}'(x)\,\frac{\partial B}{\partial \hat{v}}\Big] = \mathbb{E}_{n,g}\Big[\mathbf{\Lambda}'(x)\Big]\,(1+\delta s) \;>\; 0. \tag{A.17}$$

The implicit function theorem implies

$$\frac{\partial w^*}{\partial \hat{v}} = -\frac{F_{\hat{v}}}{F_w} > 0, \qquad \text{since } F_{\hat{v}} > 0 \text{ and } F_w < 0. \tag{A.18}$$

Higher skill raises expected surplus B, making the hire more valuable. This creates incentive to raise wages to improve acceptance probability. However, this operates through a surplus trade-off rather than a mechanical acceptance effect: the firm balances higher surplus against the marginal cost of wage adjustment (including the deviation cost κ_w).

Wage rigidity κ_w. From the implicit-function formulas,

$$\frac{\partial w^*}{\partial \hat{v}} = -\frac{F_{\hat{v}}}{F_w}, \qquad \frac{\partial w^*}{\partial \hat{\omega}} = -\frac{F_{\hat{\omega}}}{F_w}, \qquad F_w < 0.$$

The key object is F_w, which can be computed as:

$$\begin{aligned} F_w &= \mathbb{E}_{n,g}\Big[\mathbf{\Lambda}''(x)B + 2\mathbf{\Lambda}'(x)\frac{\partial B}{\partial w}\Big] - \mathbb{E}_{n,g}\Big[\mathbf{\Lambda}(x)\Big]\,\kappa_w \\ &= \underbrace{\mathbb{E}_{n,g}\Big[\mathbf{\Lambda}''(x)B\Big]}_{\text{curvature term}} - \underbrace{2\,\mathbb{E}_{n,g}\Big[\mathbf{\Lambda}'(x)\Big]\,[1+\kappa_w(w-w_g^{\text{eff}})]}_{\text{marginal-cost term}} - \underbrace{\kappa_w\,\mathbb{E}_{n,g}\Big[\mathbf{\Lambda}(x)\Big]}_{\text{direct rigidity term}}. \end{aligned} \tag{A.19}$$

Because $\mathbb{E}_{n,g}[\Lambda(x)] > 0$, a higher κ_w makes F_w more negative by approximately $\mathbb{E}_{n,g}[\Lambda(x)]$ per unit increase in κ_w. Therefore $|F_w|$ increases with κ_w, and both $|\partial w^*/\partial \hat{v}|$ and $|\partial w^*/\partial \hat{\omega}|$ shrink. Thus wage rigidity dampens both channels.

The model delivers conditions under which wage-setting is more sensitive to outside-option beliefs than to skill beliefs.

Using the implicit-function representation

$$\frac{\partial w^*}{\partial \theta} = -\frac{F_\theta}{F_w}, \qquad F_w < 0,$$

relative sensitivity is governed by the numerators $F_{\hat{\omega}}$ and $F_{\hat{v}}$. Outside-option effects are more likely to dominate whenever

$$|F_{\hat{\omega}}\,\Delta\hat{\omega}| > |F_{\hat{v}}\,\Delta\hat{v}|.$$

From the earlier derivations,

$$F_{\hat{\omega}} = \mathbb{E}_{n,g}\Big[-\Lambda''(x)\,B(v,w,s)\Big] + \mathbb{E}_{n,g}\Big[\Lambda'(x)\Big]\,[1+\kappa_w(w-w_g^{\text{eff}})] + \alpha\,\kappa_w\,\mathbb{E}_{n,g}\Big[\Lambda(x)\Big],$$
$$F_{\hat{v}} = \mathbb{E}_{n,g}\Big[\Lambda'(x)\Big]\,(1+\delta s).$$

The model implies that outside-option dominance is more likely under the following conditions:

(i) **Failed hires are costly.** A larger rejection cost c_R raises the value of closing, $B = S + c_R$, which scales the curvature-weighted term $\mathbb{E}[-\Lambda''(x)B]$ and amplifies the marginal value of increasing acceptance.

(ii) **Offers lie on the high-acceptance side of the interior.** For the logistic function,

$$\Lambda''(x) = \Lambda'(x)\Big(1 - 2\Lambda(x)\Big).$$

Hence $-\Lambda''(x) > 0$ exactly when $\Lambda(x) > \frac{1}{2}$. When the posterior places substantial mass in this region (acceptance is more likely than not but not near certainty), the term $\mathbb{E}[-\Lambda''(x)B]$ is positive and increasing in B, strengthening the outside-option (closing) channel.

Under these conditions, wages function primarily as a closing instrument, responding strongly to inferred outside options, while skill assessments are more naturally absorbed by less rigid margins such as senior-track recommendations.

A.3.2 Senior-Track Recommendation and Responses

Senior FOC (full expectation). Differentiate (A.8) w.r.t. s. Now $\partial x/\partial s = \tau$ and $\partial B/\partial s = \partial S/\partial s = \delta v - \kappa_s(s - \bar{s})$. The senior FOC is

$$0 = \frac{\partial \Pi}{\partial s} = \mathbb{E}_{n,g}\Big[\tau\, \Lambda'(x)\, B(v, w, s) + \Lambda(x)\, \Big(\delta v - \kappa_s(s - \bar{s})\Big)\Big] \tag{A.20}$$

Total differentiation for comparative statics. Fix w at the wage best response or treat (w, s) jointly; for sign arguments it is enough to consider the implicit equation in s. Define

$$H(s; n, g, w) \equiv \mathbb{E}_{n,g}\Big[\tau\, \Lambda'(x)\, B(v, w, s) + \Lambda(x)\, \Big(\delta v - \kappa_s(s - \bar{s})\Big)\Big], \quad x = w + \tau s - \omega.$$

An interior optimum satisfies $H(s^*; \cdot) = 0$ and assume local concavity $H_s(s^*; \cdot) < 0$. Then for any belief parameter $\theta \in \{\hat{v}, \hat{\omega}\}$,

$$\frac{\partial s^*}{\partial \theta} = -\frac{H_\theta}{H_s}, \qquad H_s < 0. \tag{A.21}$$

Skill effect on s^*. Treat $v = \hat{v} + \tilde{v}$ with fixed posterior shape. Holding s fixed, x does not depend on v, hence $x_{\hat{v}} = 0$. Compute $H_{\hat{v}}$ by differentiating inside the expectation:

$$\begin{aligned} H_{\hat{v}} &= \mathbb{E}_{n,g}\Big[\tau\, \Lambda'(x)\, \frac{\partial B}{\partial \hat{v}} + \Lambda(x)\, \frac{\partial}{\partial \hat{v}}(\delta v)\Big] \\ &= \mathbb{E}_{n,g}\Big[\tau\, \Lambda'(x)\, (1 + \delta s) + \Lambda(x)\, \delta\Big] \;>\; 0. \end{aligned} \tag{A.22}$$

Therefore, since $H_s < 0$,

$$\frac{\partial s^*}{\partial \hat{v}} = -\frac{H_{\hat{v}}}{H_s} > 0 \tag{A.23}$$

Higher inferred skill raises the marginal productivity return from senior placement (δv) and also raises surplus, making acceptance more valuable.

Outside-option effect on s^*. Treat $\omega = \hat{\omega} + \tilde{\omega}$, so $x_{\hat{\omega}} = -1$ and $B_{\hat{\omega}} = 0$. Then

$$\begin{aligned} H_{\hat{\omega}} &= \mathbb{E}_{n,g}\Big[\tau\,\Lambda''(x)\,x_{\hat{\omega}}\,B(v,w,s) + \Lambda'(x)\,x_{\hat{\omega}}\,\Big(\delta v - \kappa_s(s-\bar{s})\Big)\Big] \\ &= \mathbb{E}_{n,g}\Big[-\tau\,\Lambda''(x)\,B(v,w,s) - \Lambda'(x)\,\Big(\delta v - \kappa_s(s-\bar{s})\Big)\Big]. \end{aligned} \tag{A.24}$$

Hence

$$\frac{\partial s^*}{\partial \hat{\omega}} = -\frac{H_{\hat{\omega}}}{H_s},$$

and the sign depends on the posterior distribution and parameter values.

Unlike wages, senior-track recommendations do not respond monotonically to outside options. The first term in $H_{\hat{\omega}}$ captures an acceptance effect (higher $\hat{\omega}$ makes acceptance harder, raising the value of instruments that improve acceptance). But the second term shows that when the productivity margin $\delta v - \kappa_s(s-\bar{s})$ is positive, the direct negative effect on acceptance (through $-\Lambda'(x)$) can dominate.

A.4 Differential Instrument Responses to Skill vs. Outside Options

We now examine whether wages and senior-track recommendations respond differentially to skill versus outside-option assessments. We show that the model allows for asymmetric responses, with certain parameters favoring stronger senior responsiveness to skill while simultaneously favoring stronger wage responsiveness to outside options.

Since $F_w < 0$ and $H_s < 0$ (concavity), the implicit function theorem gives:

$$\frac{\partial s^*}{\partial \hat{v}} > \frac{\partial w^*}{\partial \hat{v}} \quad \Longleftrightarrow \quad \frac{H_{\hat{v}}}{|H_s|} > \frac{F_{\hat{v}}}{|F_w|},$$

and similarly for outside options:

$$\frac{\partial s^*}{\partial \hat{\omega}} > \frac{\partial w^*}{\partial \hat{\omega}} \quad \Longleftrightarrow \quad \frac{H_{\hat{\omega}}}{|H_s|} > \frac{F_{\hat{\omega}}}{|F_w|}.$$

Each comparison involves numerators (measuring how first-order conditions respond to belief changes) and denominators (measuring curvature of the objective function). We identify factors that can generate differential instrument responses.

A.4.1 Skill Responses

From earlier derivations:

$$F_{\hat{v}} = \mathbb{E}\Big[\Lambda'(x)\Big](1+\delta s)$$
$$H_{\hat{v}} = \mathbb{E}\Big[\tau\Lambda'(x)(1+\delta s)+\Lambda(x)\delta\Big]$$

When candidates value senior placement ($\tau > 0$) and productivity complementarity is positive ($\delta > 0$):

$$H_{\hat{v}} > F_{\hat{v}}.$$

This numerator advantage arises because senior-track recommendations benefit from both the acceptance channel (scaled by τ) and the direct productivity channel ($\Lambda(x)\delta$), while wages benefit only from the acceptance channel. Larger δ strengthens this advantage.

The second-order conditions are:

$$F_w = \mathbb{E}\Big[\Lambda''(x)[S+c_R] - 2\Lambda'(x)[1+\kappa_w(w-w_g^{\text{eff}})] - \Lambda(x)\kappa_w\Big]$$
$$H_s = \mathbb{E}\Big[\tau^2\Lambda''(x)[S+c_R] + 2\tau\Lambda'(x)[\delta v - \kappa_s(s-\bar{s})] - \Lambda(x)\kappa_s\Big]$$

The relative magnitudes depend on: (i) institutional rigidity (κ_w vs κ_s), (ii) candidate valuation of senior placement (τ), and (iii) productivity complementarity (δ).

In settings with high acceptance rates–such as our empirical context where mean acceptance is approximately 0.9–the logistic derivatives become small. For example:

$$\Lambda(x) \approx 0.9, \quad \Lambda'(x) \approx 0.09, \quad |\Lambda''(x)| \approx 0.07.$$

When $\Lambda(x)$ dominates its derivatives by an order of magnitude, the rigidity terms $-\Lambda(x)\kappa_w$ and $-\Lambda(x)\kappa_s$ dominate acceptance curvature and cross-derivative effects, giving:

$$\frac{|H_s|}{|F_w|} \approx \frac{\kappa_s}{\kappa_w}.$$

This ratio compares raw curvatures, but w and s are measured on different natural scales—wages vary continuously in dollars, while s is bounded on $[0,1]$—so it is not yet the economically relevant comparison. Writing $\tilde{w} = w/\sigma_w$, $\tilde{s} = s/\sigma_s$ for the standard deviations of w, s, the chain rule gives $\partial^2\Pi/\partial\tilde{s}^2 = \sigma_s^2 H_s$ and $\partial^2\Pi/\partial\tilde{w}^2 = \sigma_w^2 F_w$, so the standardized version of the ratio

above is

$$\frac{|\tilde{H}_{\tilde{s}}|}{|\tilde{F}_{\tilde{w}}|} \approx \frac{\sigma_s^2 \kappa_s}{\sigma_w^2 \kappa_w}.$$

The same rescaling strengthens the numerator condition above from $\tau > 0$ to $\tau > \sigma_s/\sigma_w$—a mild condition, satisfied whenever candidates place non-negligible value on advancement relative to the wage-seniority scale ratio. We evaluate both standardized conditions against the structural estimates in Appendix C.

If wage-setting faces stronger institutional constraints than senior-track placement decisions, in comparable units—formally, $\kappa_w \sigma_w^2 > \kappa_s \sigma_s^2$—the denominator ratio favors senior responsiveness.

Factors favoring stronger senior skill responsiveness. Combining numerator and denominator effects, senior-track recommendations respond more strongly to skill when:

- *Numerator*: Candidates value advancement (τ large) and productivity complementarity is strong ($\delta > 0$), ensuring $H_{\hat{v}} > F_{\hat{v}}$
- *Denominator*: Wage-setting is more rigid than senior-track placement, in comparable units, ensuring $|H_s| < |F_w|$ under high acceptance

When these factors align, both numerator and denominator favor stronger senior responsiveness: $\partial s^*/\partial \hat{v} > \partial w^*/\partial \hat{v}$.

A.4.2 Outside-Option Responses

From earlier derivations:

$$F_{\hat{\omega}} = \mathbb{E}\Big[-\Lambda''(x) B(v,w,s) \Big] + \mathbb{E}\Big[\Lambda'(x)\Big][1 + \kappa_w (w - w_g^{\text{eff}})] + \alpha \kappa_w \mathbb{E}\Big[\Lambda(x)\Big]$$
$$H_{\hat{\omega}} = \mathbb{E}\Big[-\tau \Lambda''(x) B(v,w,s) - \Lambda'(x)(\delta v - \kappa_s (s - \bar{s})) \Big]$$

where $B = S + c_R$ is the net benefit of acceptance.

For wages, all of the terms are positive when acceptance is interior ($\Lambda > 1/2$ implies $\Lambda'' < 0$):

- Acceptance curvature: $\mathbb{E}[-\Lambda'' B] > 0$
- Direct closing cost: $\mathbb{E}[\Lambda'][1 + \kappa_w (w - w_g^{\text{eff}})] > 0$

- Benchmark shift: $\alpha \kappa_w \mathbb{E}[\Lambda(x)] > 0$ unconditionally for $\alpha > 0$

Therefore $F_{\hat{\omega}} > 0$, and wages respond positively to outside options.

For senior-track recommendations, two forces offset:

- Acceptance channel: $\mathbb{E}[-\tau \Lambda'' B] > 0$ (using s to enhance acceptance, scaled by τ)
- Productivity dilution: $-\mathbb{E}[\Lambda'(\delta v - \kappa_s(s - \bar{s}))]$ is negative when productivity margin is positive

The second term reflects that higher outside options reduce acceptance probability, which reduces the expected productivity return from senior placement (output δv is realized only upon acceptance). When productivity complementarity is strong, this dilution effect can dominate, making $H_{\hat{\omega}}$ small or even negative.

The model cannot analytically determine which instrument responds more strongly to outside options–it requires empirical validation.

A.4.3 Differential Responses Through Complementarity

Role of productivity complementarity. Notably, larger δ has *opposing effects* across the two comparisons:

- *Skill responsiveness*: Larger δ increases $H_{\hat{v}}$ through the direct productivity channel, making senior more responsive to skill assessments
- *Outside-option responsiveness*: Larger δ increases the negative productivity dilution term in $H_{\hat{\omega}}$, making senior less responsive (or even negatively responsive) to outside-option assessments

This creates the potential for differential instrument responses: the same parameter (δ large) that makes senior respond more strongly to skill assessments simultaneously makes wages respond more strongly to outside-option assessments.

This complementarity-driven pattern of differential responses is reinforced when: (i) candidates value advancement (τ large), (ii) wage-setting is more rigid than senior-track placement, in comparable units, and (iii) acceptance rates are high (making rigidity dominate curvature).

The key insight is that both instruments respond to both signals, but parameters can create asymmetry in the magnitudes. When productivity complementarity is strong, we expect

senior-track recommendations to be relatively more responsive to skill signals and wages to be relatively more responsive to outside-option signals.

A.5 Group Differences and Negotiation Intensity

Using the mean-slope formulas (A.5)–(A.6), any outcome $y(w^*, s^*)$ responds to n via the chain rule. In particular, for each group g,

$$\frac{dw_g^*}{dn} = \frac{\partial w^*}{\partial \hat{\omega}} \frac{d\hat{\omega}_g}{dn} + \frac{\partial w^*}{\partial \hat{v}} \frac{d\hat{v}_g}{dn}, \tag{A.25}$$

$$\frac{ds_g^*}{dn} = \frac{\partial s^*}{\partial \hat{\omega}} \frac{d\hat{\omega}_g}{dn} + \frac{\partial s^*}{\partial \hat{v}} \frac{d\hat{v}_g}{dn}. \tag{A.26}$$

Decomposition of group differences. The difference in wage responses across groups is

$$\frac{dw_m^*}{dn} - \frac{dw_f^*}{dn} = \frac{\partial w^*}{\partial \hat{\omega}} \underbrace{\left(\frac{d\hat{\omega}_m}{dn} - \frac{d\hat{\omega}_f}{dn} \right)}_{\text{outside-option updating gap}} + \frac{\partial w^*}{\partial \hat{v}} \underbrace{\left(\frac{d\hat{v}_m}{dn} - \frac{d\hat{v}_f}{dn} \right)}_{\text{skill updating gap}}. \tag{A.27}$$

Similarly for senior-track recommendations.

Subtract the wage-gap from the senior-gap and use the chain-rule decompositions

$$\frac{ds_m^*}{dn} - \frac{ds_f^*}{dn} - \left(\frac{dw_m^*}{dn} - \frac{dw_f^*}{dn} \right) = \left(\frac{\partial s^*}{\partial \hat{\omega}} - \frac{\partial w^*}{\partial \hat{\omega}} \right) \Delta_\omega + \left(\frac{\partial s^*}{\partial \hat{v}} - \frac{\partial w^*}{\partial \hat{v}} \right) \Delta_v$$

Rearranging, we obtain

$$\left(\frac{\partial s^*}{\partial \hat{v}} - \frac{\partial w^*}{\partial \hat{v}} \right) \Delta_v > \left(\frac{\partial w^*}{\partial \hat{\omega}} - \frac{\partial s^*}{\partial \hat{\omega}} \right) \Delta_\omega$$

A.6 Firm Surplus and the Rejection Threshold

The effect of belief updates on expected profit at the optimum (w^*, s^*) can be signed directly.

Effect of perceived productivity. Since output is $v(1 + \delta s)$ and is realized upon acceptance:

$$\frac{\partial \Pi^*}{\partial \hat{v}} = \mathbb{E}_{n,g} \left[\Lambda(x)(1 + \delta s^*) \right] > 0. \tag{A.28}$$

Higher perceived productivity raises expected output and therefore expected surplus. This effect is unambiguously positive.

Effect of perceived outside options. Higher outside options reduce acceptance probability for any fixed wage. At the optimum:

$$\frac{\partial \Pi^*}{\partial \hat{\omega}} = -\mathbb{E}_{n,g}\left[\Lambda'(x) \cdot B(v, w^*, s^*)\right] < 0, \tag{A.29}$$

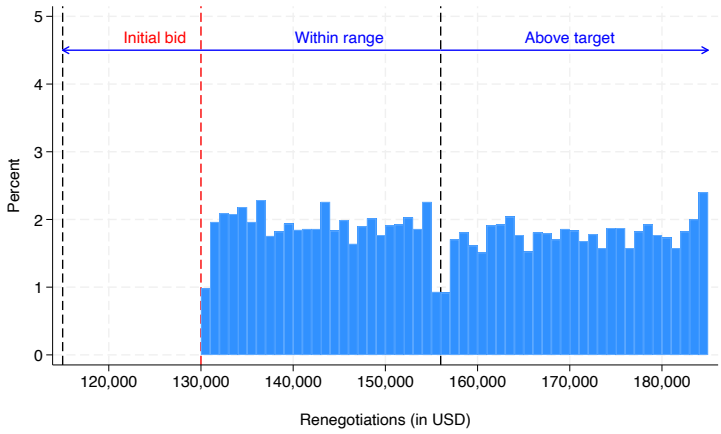
where $B = S + c_R > 0$ is the net value of closing the hire. Since $\Lambda'(x) > 0$ and we expect $B > 0$, this derivative tends to be negative. The firm optimally raises wages in response to higher $\hat{\omega}$, but this adjustment is costly leaving expected profit lower.

Implication for rejection. Since $\partial \Pi^*/\partial \hat{\omega} < 0$ and $\partial \Pi^*/\partial \hat{v} > 0$, expected profit is decreasing in perceived outside options and increasing in perceived productivity. The firm's participation constraint $\Pi^* \geq \bar{V}$ is therefore most likely to bind — and the candidacy to be rejected — when high counteroffers raise perceived outside options without a corresponding upward revision in perceived productivity.

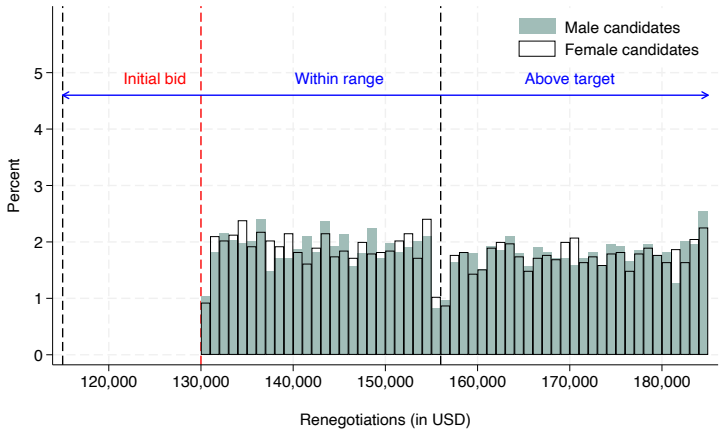
B Additional Results

FIGURE B1. Distribution of Counteroffers

(A) Whole Sample

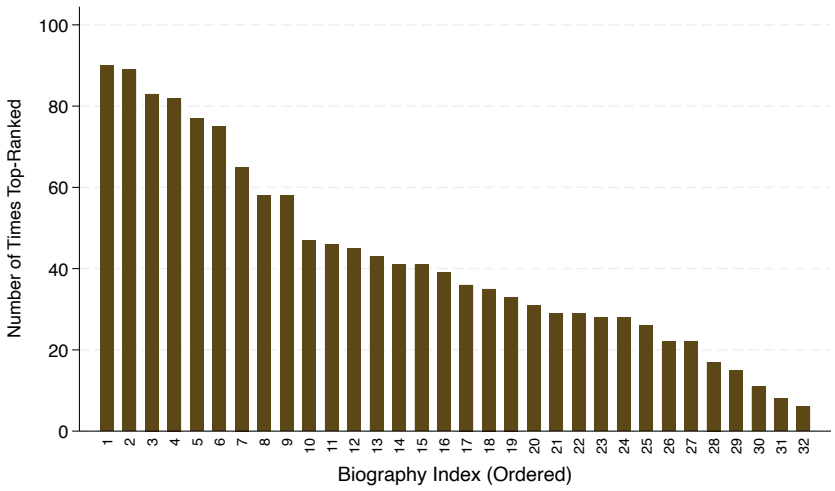


(B) By Gender



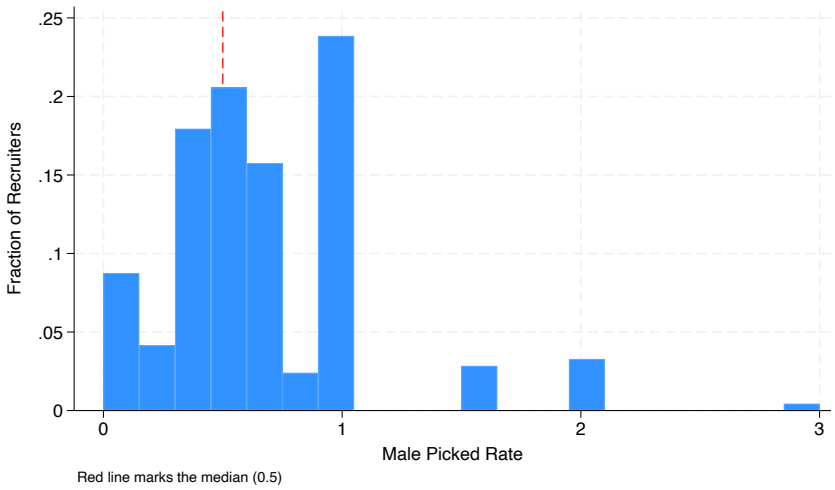
Notes: This figure shows the distribution of hypothetical counteroffers made by fictitious candidates for the whole sample (Panel A) and separately by gender (Panel B). Attached to each candidate's resume, we provided counteroffer salaries to the firm's initial offer of $130,000. These counteroffers, all above the initial offers, were randomly assigned and represent negotiation intensity. A certain portion of the candidates were shown to present counteroffers below the maximum salary range of $155,000 (but above the initial offer), while others presented counteroffers above the maximum, up to 20 percent above the threshold (i.e., from $156,000 to $185,000).

FIGURE B2. Distribution of Biographies Mentioned as Top by Recruiters



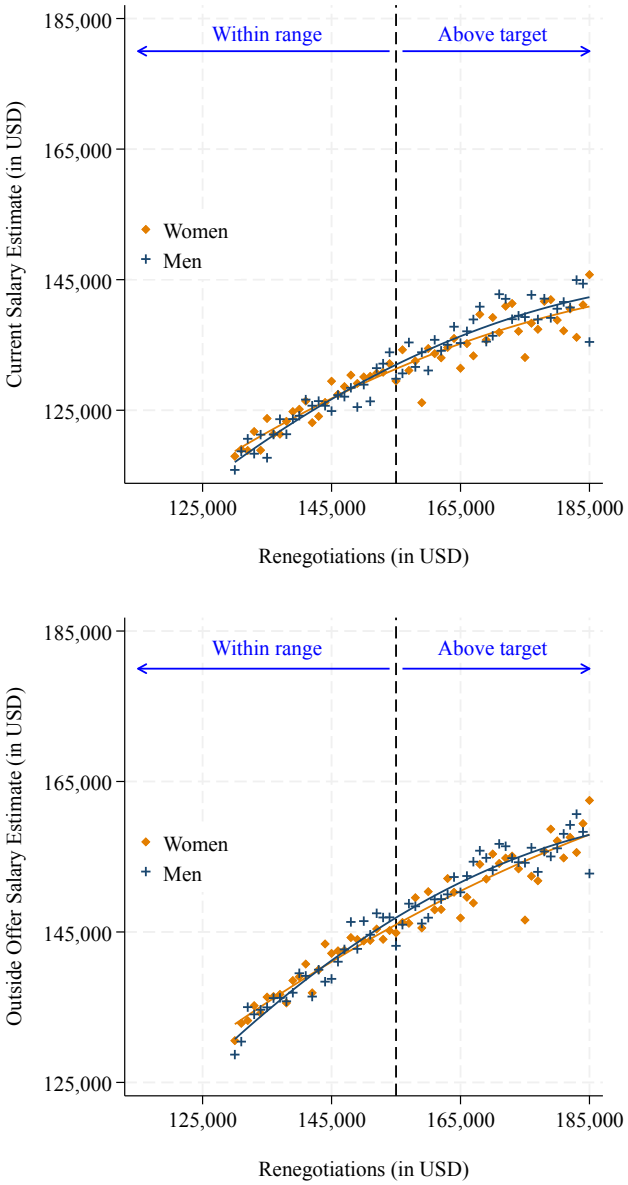
Notes: This figure shows the distribution of biographies that were mentioned as one of the top three by recruiters. It illustrates how many times each biography was mentioned, regardless of the assigned gender. There are 32 biographies in total. The number of times each biography was presented to recruiters ranges from 231 to 252, with a mean of 242.

FIGURE B3. Distribution of Male-Favoring Propensity Across Recruiters



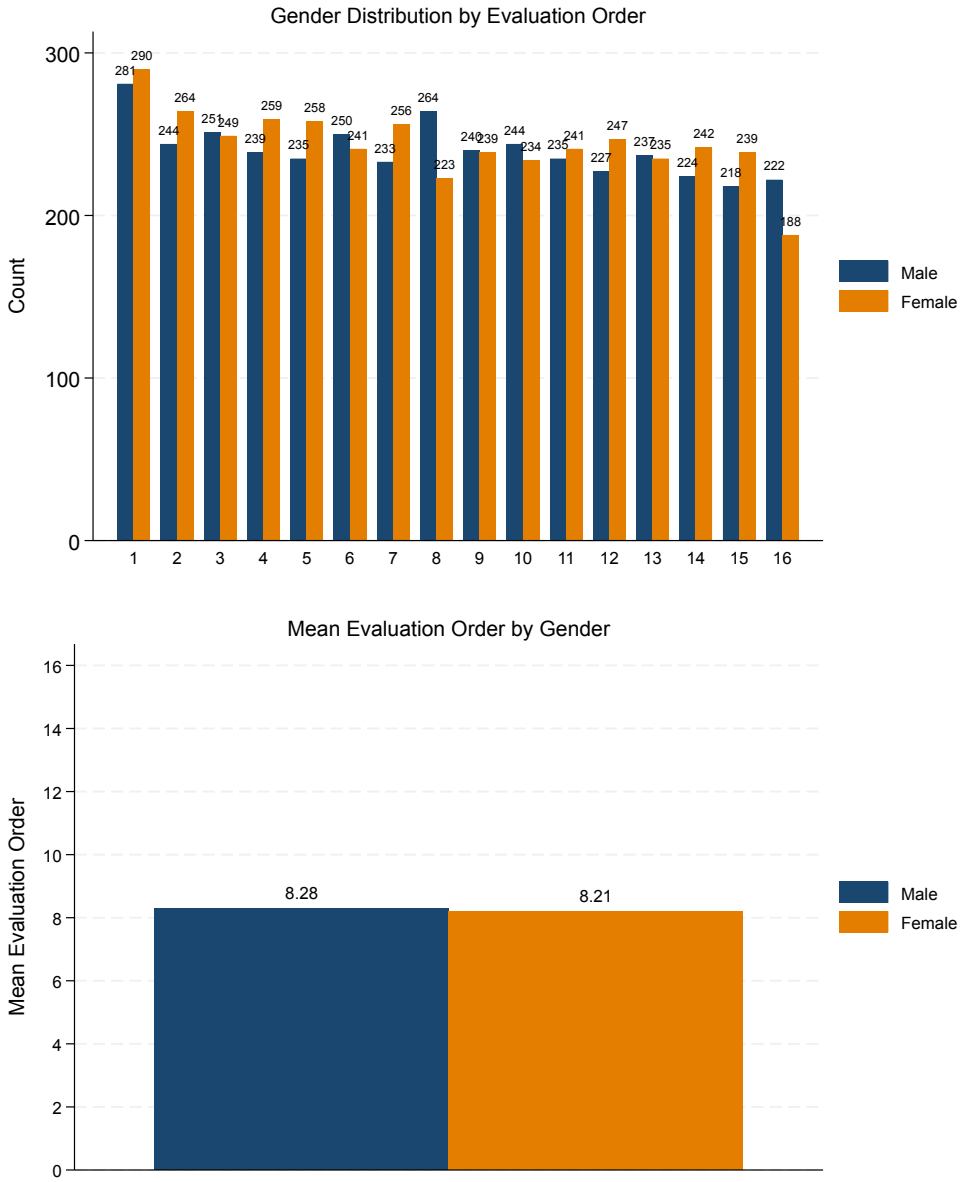
Notes: We construct a male-favoring propensity score for each recruiter. We use a pre-registered method to assign male-favoring propensities to recruiters. From the post-experimental survey, we asked each recruiter to name the top three candidates. We then compile a list of the top 6 biographies mentioned by all recruiters and calculate for each recruiter the ratio of male resumes selected to the total number of top 6 resumes presented to them with male names. Since it is possible that a recruiter may never be shown a top male candidate by chance, we add one to the denominator to avoid division by zero. If a recruiter did not rank anyone, we treat this as missing data for their propensity.

FIGURE B4. Current Salary and Outside Option Estimates from Higher Counteroffers



Notes: This figure illustrates the relationship between hypothetical counteroffers made by fictitious candidates and various perceptions and decisions by recruiters, separately by gender. The figure is organized clockwise from the top: current salary estimates and best outside options. To create the graphs, we calculated the mean of the outcome variable within each $1,000 increment of the counteroffer range.

FIGURE B5. Evaluation Order



Notes: The top panel displays the count of male and female resumes evaluated at each position in the sequence, showing how gender was distributed across evaluation order. The bottom panel presents the mean outcome variable by evaluation position, averaged separately for male and female candidates

Table B1 – Balancing Test Over Salary Range - Recruiters

Variables	Control	Treatment	Difference	p-value
Sociodemographic characteristics				
Female	0.738	0.732	-0.006	0.711
White	0.536	0.535	-0.000	0.994
Black	0.250	0.240	-0.009	0.561
Total Hours	390.220	355.803	-34.417	0.477
Hourly Rate	63.738	64.555	0.818	0.483
Bachelor degree	0.892	0.900	0.008	0.473
5 years of experience or more	0.766	0.765	-0.001	0.928
Test of joint significance	F-stat: 0.300 (p-value: 0.945)			
Resume characteristics				
Female CV	0.508	0.500	-0.008	0.457
Total Experince (Month)	54.009	53.633	-0.377	0.369
N	3,790	3,959		

Notes: We perform data quality checks to ensure that the randomization was successful. In this table, we test whether resumes with counteroffers exceeding the maximum salary range were more likely to be evaluated by certain types of recruiters or had different resume characteristics. We compare the types of recruiters (e.g., experience, race, and gender) evaluating counteroffers made below or above the target range, and resume characteristics that are considered under the two conditions.

Table B2 – Balancing Test Over Pay Transparency - Recruiters

Variables	Control	Treatment	Difference	p-value
Sociodemographic characteristics				
Female	0.733	0.737	0.003	0.934
White	0.530	0.541	0.011	0.815
Black	0.259	0.232	-0.027	0.490
Total Hours	315.321	426.133	110.812	0.335
Hourly Rate	64.960	63.405	-1.555	0.577
Bachelor degree	0.892	0.900	0.008	0.760
5 years of experience or more	0.809	0.725	-0.085	0.027
Test of joint significance	F-stat: 1.100 (p-value: 0.349)			
Resume characteristics				
Female CV	0.503	0.504	0.001	0.671
Total Experince (Month)	54.072	53.579	-0.493	0.277
Counter offer over target	0.503	0.518	0.015	0.435
Percent over target: 25	0.342	0.301	-0.040	0.340
Percent over target: 50	0.303	0.327	0.024	0.571
Percent over target: 75	0.355	0.371	0.017	0.705
N	3,741	4,008		

Notes: We perform a balance test by regressing recruiters' characteristics on indicators for the pay transparency treatment.

Table B3 – Attrition

	Probability of Not Completing the Task				
Recruiters' characteristics:					
Female	-0.002				
	(0.012)				
White		0.008			
		(0.010)			
Black			0.006		
			(0.013)		
Bachelor				-0.032	
				(0.026)	
Work Experience: 5 years or more					0.004
					(0.011)
Y mean	.017	.017	.017	.017	.017
N	7,749	7,749	7,749	7,749	7,749

Notes: This table examines how attrition varies with recruiters' sociodemographic characteristics. Each column reports the estimate of a separate univariate regression of the covariate on recruiters' probability of completing less than 16 candidates' evaluations. Standard errors are clustered at the recruiter level.

Table B4 – Effect of Female Candidate CV

Variables	Male Candidate (1)	Female Candidate (2)	Difference (2)-(1)	p-val > ue
Recommendations				
Salary offer	145,880	145,400	-433.14*	0.090
			(255.18)	
Accept counter offer	0.388	0.408	0.022**	0.032
			(0.010)	
Reject application	0.102	0.096	-0.006	0.331
			(0.006)	
Recommends WFH offer	0.598	0.633	0.038***	0.000
			(0.009)	
Recommend for senior position	0.616	0.619	0.006	0.554
			(0.010)	
Market Value and Options				
Current salary estimate	131,816	131,099	-652.49*	0.051
			(332.87)	
Best outside offer estimate	146,695	146,100	-501.12	0.104
			(308.00)	
Candidate will likely accept	0.809	0.804	-0.005	0.522
			(0.008)	
Skill Assessments				
Technical skills - High	0.809	0.813	0.005	0.517
			(0.007)	
Team skills - High	0.805	0.819	0.016**	0.035
			(0.007)	
Initiative skills - High	0.819	0.839	0.021***	0.003
			(0.007)	
N			7,749	
Number of clusters			493	

Notes: This table shows how recruiters perceive and evaluate male and female candidates across various dimensions. The first section presents recommendations on the final salary offer, whether to accept the candidate's counteroffer, reject the application altogether, offer work-from-home arrangements, and recommend the candidate for a new senior-track position. The second section reports recruiters' estimates of the candidate's market value and likelihood of accepting the offer. The third section provides results for recruiters' assessments of candidates' skills. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B5 – Robustness Checks - Effect of Counteroffer Over Salary Range

Variables	Baseline		With Controls		With Bio FE		With Recruiter FE		With Order FE	
	Difference	*p*-value	Difference	*p*-value	Difference	*p*-value	Difference	*p*-value	Difference	*p*-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Recommendations										
Salary offer	14,591.21***	0.000	14,575.53***	0.000	14,567.15***	0.000	14,043.87***	0.000	14,584.31***	0.000
	(396.32)		(393.67)		(390.60)		(370.45)		(396.68)	
Accept counter offer	-0.309***	0.000	-0.309***	0.000	-0.309***	0.000	-0.339***	0.000	-0.308***	0.000
	(0.015)		(0.015)		(0.015)		(0.014)		(0.015)	
Reject application	0.044***	0.000	0.044***	0.000	0.046***	0.000	0.053***	0.000	0.044***	0.000
	(0.008)		(0.008)		(0.008)		(0.008)		(0.008)	
Recommends WFH offer	0.058***	0.000	0.058***	0.000	0.058***	0.000	0.067***	0.000	0.058***	0.000
	(0.016)		(0.015)		(0.016)		(0.011)		(0.016)	
Recommend for senior position	0.088***	0.000	0.089***	0.000	0.089***	0.000	0.094***	0.000	0.088***	0.000
	(0.014)		(0.014)		(0.013)		(0.012)		(0.014)	
Market Value and Options										
Current salary estimate	12,170.52***	0.000	12,163.75***	0.000	12,152.54***	0.000	13,198.11***	0.000	12,138.33***	0.000
	(670.58)		(656.49)		(654.74)		(497.31)		(672.05)	
Best outside offer estimate	13,301.13***	0.000	13,251.93***	0.000	13,249.12***	0.000	13,238.70***	0.000	13,281.85***	0.000
	(578.18)		(558.98)		(553.14)		(441.28)		(578.77)	
Candidate will likely accept	-0.170***	0.000	-0.170***	0.000	-0.171***	0.000	-0.199***	0.000	-0.169***	0.000
	(0.012)		(0.012)		(0.012)		(0.012)		(0.012)	
Skill Assessments										
Technical skills - High	0.021*	0.065	0.022*	0.055	0.020*	0.072	0.020**	0.019	0.021*	0.061
	(0.011)		(0.011)		(0.011)		(0.009)		(0.011)	
Team skills - High	0.012	0.272	0.012	0.274	0.011	0.321	0.002	0.821	0.012	0.277
	(0.011)		(0.011)		(0.011)		(0.008)		(0.011)	
Initiative skills - High	0.031***	0.005	0.031***	0.003	0.030***	0.004	0.024***	0.003	0.031***	0.005
	(0.011)		(0.011)		(0.011)		(0.008)		(0.011)	
N	7,749		7,749		7,749		7,745		7,749	
Number of clusters	493		493		493		493		493	
Controls	No		Yes		Yes		No		No	
Bio FE	No		No		Yes		No		No	
Recruiters FE	No		No		No		Yes		No	
Order FE	No		No		No		No		Yes	

Notes: This table provides robustness checks for Table 2, presenting results when including only specific subsets of controls. For variable descriptions, refer to the notes in Table 2. Column (1): baseline with no controls. Column (2): adds recruiter-level controls. These controls include recruiter gender, race, total platform hours, hourly rate, a bachelor's degree indicator, and a 5+ years of recruiting experience indicator. Column (3): adds biographical fixed effects, absorbing the identity of each unique resume template. Column (4): recruiter fixed effects, absorbing all recruiter-level heterogeneity. Column (5): evaluation order fixed effects, where the order is determined by the chronological sequence in which each recruiter submitted their evaluations. Standard errors clustered at the recruiter level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B6 – Robustness Checks - Effect of Counteroffer Over salary Range By Gender

Variables	Baseline > 6			With Controls			With Bio FE			With Recruiter FE			> With Order FE		
	Female CV	**Male CV**	**Difference**	**Female CV**	**Male CV**	**Difference**	**Female CV**	**Male CV**	**Difference**	**Female CV**	**Male CV**	**Difference**	**Female CV**	**Male CV**	**Difference**
Recommendations															
Salary offer	14,353.49***	14,827.11***	-473.62	14,333.81***	14,808.60***	-465.23	14,377.57***	14,750.51***	-379.02	14,028.03***	14,183.18***	-505.82	14,584.31***	14,584.31***	-427.65
	(484.28)	(425.38)	(449.48)	(479.26)	(423.78)	(444.27)	(468.18)	(421.52)	(428.08)	(462.49)	(440.57)	(413.96)	(396.68)	(396.68)	(451.88)
Accept counter offer	-0.315***	-0.302***	-0.014	-0.315***	-0.302***	-0.013	-0.315***	-0.305***	-0.009	-0.340***	-0.342***	-0.007	-0.308***	-0.308***	-0.013
	(0.018)	(0.018)	(0.021)	(0.018)	(0.018)	(0.021)	(0.018)	(0.018)	(0.021)	(0.018)	(0.018)	(0.021)	(0.015)	(0.015)	(0.021)
Reject application	0.059***	0.029***	0.029**	0.059***	0.030***	0.029**	0.059***	0.033***	0.027**	0.068***	0.046***	0.027**	0.044***	0.044***	0.029**
	(0.010)	(0.011)	(0.014)	(0.010)	(0.011)	(0.014)	(0.010)	(0.011)	(0.013)	(0.011)	(0.011)	(0.013)	(0.008)	(0.008)	(0.014)
Recommends WFH offer	0.066***	0.050**	0.016	0.067***	0.049**	0.017	0.066***	0.048**	0.017	0.067***	0.061***	0.008	0.058***	0.058***	0.017
	(0.018)	(0.020)	(0.021)	(0.018)	(0.019)	(0.021)	(0.018)	(0.019)	(0.021)	(0.014)	(0.014)	(0.017)	(0.016)	(0.016)	(0.021)
Recommend for senior position	0.058***	0.120***	-0.062***	0.058***	0.121***	-0.063***	0.059***	0.117***	-0.059***	0.052***	0.127***	-0.068***	0.088***	0.088***	-0.062***
	(0.018)	(0.017)	(0.022)	(0.018)	(0.017)	(0.022)	(0.017)	(0.016)	(0.021)	(0.018)	(0.016)	(0.021)	(0.014)	(0.014)	(0.022)
Market Value and Options															
Current salary estimate	11,015.66***	13,334.72***	-2,319.07**	11,052.74***	13,266.39***	-2,212.36**	11,154.16***	13,175.99***	-2,099.97**	13,098.74***	13,492.66***	-1,209.12*	12,138.33***	12,138.33***	-2,270.59**
	(920.45)	(709.41)	(950.98)	(902.75)	(686.82)	(924.15)	(880.85)	(680.67)	(905.40)	(640.57)	(591.09)	(641.96)	(672.05)	(672.05)	(951.17)
Best outside offer estimate	12,681.56***	13,923.24***	-1,241.68*	12,634.68***	13,851.85***	-1,205.74*	12,734.99***	13,757.60***	-1,074.23	13,195.29***	13,361.41***	-898.28	13,281.85***	13,281.85***	-1,185.89
	(655.94)	(714.28)	(734.66)	(637.57)	(689.23)	(715.75)	(616.88)	(675.63)	(683.45)	(569.63)	(571.68)	(600.73)	(578.77)	(578.77)	(738.38)
Candidate will likely accept	-0.183***	-0.158***	-0.025	-0.183***	-0.158***	-0.025	-0.182***	-0.159***	-0.023	-0.215***	-0.180***	-0.027	-0.169***	-0.169***	-0.024
	(0.016)	(0.015)	(0.017)	(0.016)	(0.015)	(0.017)	(0.016)	(0.015)	(0.017)	(0.015)	(0.014)	(0.016)	(0.012)	(0.012)	(0.017)
Skill Assessments															
Technical skills - High	0.005	0.037**	-0.032*	0.006	0.037**	-0.032*	0.005	0.034**	-0.030	0.007	0.029**	-0.034**	0.021	0.021	-0.030
	(0.014)	(0.015)	(0.019)	(0.014)	(0.015)	(0.019)	(0.014)	(0.015)	(0.018)	(0.013)	(0.012)	(0.016)	(0.011)	(0.011)	(0.019)
Team skills - High	-0.008	0.033**	-0.041**	-0.009	0.033**	-0.042**	-0.008	0.031**	-0.039**	-0.017	0.025	-0.040**	0.012	0.012	-0.040**
	(0.014)	(0.015)	(0.019)	(0.014)	(0.015)	(0.019)	(0.014)	(0.015)	(0.018)	(0.013)	(0.013)	(0.017)	(0.011)	(0.011)	(0.019)
Initiative skills - High	0.015	0.047***	-0.031*	0.016	0.047***	-0.032*	0.016	0.045***	-0.029*	0.015	0.031**	-0.032**	0.031***	0.031***	-0.031*
	(0.013)	(0.015)	(0.018)	(0.013)	(0.015)	(0.018)	(0.013)	(0.014)	(0.017)	(0.012)	(0.012)	(0.015)	(0.011)	(0.011)	(0.018)
N	3,905	3,844	7,749	3,905	3,844	7,749	3,905	3,844	7,749	3,903	3,839	7,745	7,749	7,749	7,749
Number of clusters	493	493	493	493	493	493	493	493	493	493	493	493	493	493	493
Controls	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Bio FE	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No	No	No	No
Recruiters FE	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No
Order FE	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes

Notes: This table provides robustness checks for Table 3, presenting results when including only specific subsets of controls. For variable descriptions, refer to the notes in Table 3. The "Difference" column reports the gender gap (female minus male) in the over-target effect. Column (1): baseline with no controls. Column (2): adds recruiter-level controls (recruiter gender, race, total platform hours, hourly rate, bachelor's degree indicator, and 5+ years recruiting experience indicator). Column (3): adds biographical fixed effects, absorbing the identity of each unique resume template. Column (4): recruiter fixed effects, absorbing all recruiter-level heterogeneity. Column (5): evaluation order fixed effects, where the order is determined by the chronological sequence in which each recruiter submitted their evaluations. Standard errors clustered at the recruiter level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B7 – Correlation between Leadership Keywords and Skill Assessment

	Technical Skills			Team Skills			Initiative Skills		
	(1) Full	(2) Female	(3) Male	(4) Full	(5) Female	(6) Male	(7) Full	(8) Female	(9) Male
Abov. Median Leadership	0.026***	0.020	0.033***	0.059***	0.058***	0.060***	0.013*	0.006	0.020
	(0.008)	(0.012)	(0.012)	(0.009)	(0.013)	(0.013)	(0.008)	(0.011)	(0.013)
Constant	0.796***	0.802***	0.791***	0.779***	0.786***	0.772***	0.822***	0.836***	0.808***
	(0.011)	(0.013)	(0.013)	(0.012)	(0.014)	(0.014)	(0.012)	(0.013)	(0.014)
Observations	7749	3905	3844	7749	3905	3844	7749	3905	3844
R-squared	0.001	0.001	0.002	0.006	0.005	0.006	0.000	0.000	0.001

Notes: Above Median Leadership is an indicator equal to one if the candidate's count of leadership- and management-related keywords exceeds the sample median. Keywords include: *led, lead, leading, leader, leadership, managed, manage, managing, manager, management, supervised, supervise, supervising, supervision, directed, directing, director, coordinated, coordinate, coordinating, coordinator, coordination, team lead, project management, strategic planning, mentored, mentor, mentoring, mentorship, hiring, oversaw, oversee, overseeing, oversight*. Standard errors clustered at the recruiter level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B8 – Correlation between Leadership Keywords and Recruiter Decisions

	Reject			Senior Recommendation			WFH Recommendation			Offer		
	(1) All	(2) Female	(3) Male	(4) All	(5) Female	(6) Male	(7) All	(8) Female	(9) Male	(10) All	(11) Female	(12) Male
Abov. Median Leadership	-0.014**	-0.020**	-0.007	0.039***	0.021	0.057***	-0.002	-0.003	-0.002	126.511	-261.825	-261.825
	(0.006)	(0.009)	(0.010)	(0.011)	(0.016)	(0.016)	(0.009)	(0.014)	(0.015)	(273.399)	(427.252)	(427.252)
Constant	0.107***	0.108***	0.106***	0.596***	0.608***	0.584***	0.617***	0.635***	0.599***	145567.956***	145547.718***	145547.718***
	(0.007)	(0.009)	(0.009)	(0.012)	(0.015)	(0.015)	(0.016)	(0.017)	(0.018)	(340.360)	(416.487)	(416.487)
Observations	7749	3905	3844	7749	3905	3844	7749	3905	3844	7749	3905	3905
R-squared	0.001	0.001	0.000	0.002	0.000	0.003	0.000	0.000	0.000	0.000	0.000	0.000

Notes: Above Median Leadership is an indicator equal to one if the candidate's count of leadership- and management-related keywords exceeds the sample median. Reject is an indicator for application rejection. Senior Recommendation is an indicator for recommending the candidate for a senior-track position. WFH Recommendation is the recruiter's work-from-home recommendation. Offer is the final salary offer in dollars. Standard errors clustered at the recruiter level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B9 – Correlation between Resume Quality Index and Skill Assessment

	Technical Skills			Team Skills			Initiative Skills		
	(1) Full	(2) Female	(3) Male	(4) Full	(5) Female	(6) Male	(7) Full	(8) Female	(9) Male
High Quality	0.034***	0.035***	0.033***	-0.000	-0.007	0.006	0.027***	0.028***	0.026**
	(0.008)	(0.012)	(0.013)	(0.008)	(0.012)	(0.013)	(0.007)	(0.010)	(0.012)
Constant	0.799***	0.800***	0.798***	0.812***	0.821***	0.803***	0.820***	0.830***	0.811***
	(0.010)	(0.011)	(0.012)	(0.011)	(0.011)	(0.012)	(0.011)	(0.012)	(0.012)
Observations	7749	3905	3844	7749	3905	3844	7749	3905	3844
R-squared	0.002	0.002	0.002	0.000	0.000	0.000	0.001	0.001	0.001

Notes: High Quality is an indicator equal to one if the candidate's resume quality index exceeds the sample median. The index is the mean of three indicators: whether the candidate holds a computer science degree, has senior-level work experience (job title containing "II" or "Sr."), and has above-median total work experience (in months). Standard errors clustered at the recruiter level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B10 – Correlation between Resume Quality Index and Recruiter Decisions

	Reject			Senior Recommendation			WFH Recommendation			Offer		
	(1) All	(2) Female	(3) Male	(4) All	(5) Female	(6) Male	(7) All	(8) Female	(9) Male	(10) All	(11) Female	(12) Male
High Quality	-0.022***	-0.026**	-0.019*	0.056***	0.064***	0.046***	0.021**	0.011	0.032**	1281.307***	1357.596***	1208.458***
	(0.007)	(0.010)	(0.010)	(0.011)	(0.016)	(0.016)	(0.010)	(0.015)	(0.016)	(286.185)	(409.356)	(437.977)
Constant	0.107***	0.105***	0.108***	0.599***	0.597***	0.600***	0.608***	0.629***	0.587***	145203.711***	144936.639***	145473.087***
	(0.007)	(0.008)	(0.008)	(0.011)	(0.013)	(0.013)	(0.016)	(0.017)	(0.017)	(318.966)	(352.676)	(369.852)
Observations	7749	3905	3844	7749	3905	3844	7749	3905	3844	7749	3905	3844
R-squared	0.001	0.002	0.001	0.003	0.004	0.002	0.000	0.000	0.001	0.002	0.002	0.002

Notes: High Quality is an indicator equal to one if the candidate's resume quality index exceeds the sample median. The index is the mean of three indicators: whether the candidate holds a computer science degree, has senior-level work experience (job title containing "II" or "Sr."), and has above-median total work experience (in months). Reject is an indicator for application rejection. Senior Recommendation is an indicator for recommending the candidate for a senior-track position. WFH Recommendation is the recruiter's work-from-home recommendation. Offer is the final salary offer in dollars. Standard errors clustered at the recruiter level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B11 – Heterogeneity by Candidates' Graduation Year and Work Experience

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
Panel A: Candidates' Graduation Year				
Young	-765.436**	-0.051*	-0.009	0.043*
	(348.065)	(0.027)	(0.015)	(0.026)
Over target	15639.644***	0.101***	0.016	0.073**
	(668.749)	(0.027)	(0.017)	(0.031)
Over target × Young	-1225.319	0.031	0.021	-0.037
	(755.203)	(0.033)	(0.021)	(0.034)
Female CV	-522.325	0.046*	-0.026	0.072***
	(398.431)	(0.025)	(0.016)	(0.024)
Female CV × Young	614.381	-0.013	0.008	-0.069**
	(469.551)	(0.032)	(0.020)	(0.030)
Over target × Female CV	-1097.584	-0.069*	0.056**	-0.051
	(860.450)	(0.036)	(0.025)	(0.037)
Over target × Female CV × Young	923.707	0.008	-0.041	0.103**
	(1051.038)	(0.044)	(0.029)	(0.045)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749
Panel B: Candidates' Work Experience				
Experienced	1447.336***	0.102***	-0.031**	0.034
	(335.409)	(0.024)	(0.014)	(0.023)
Over target	13718.341***	0.111***	0.035**	0.053**
	(557.110)	(0.025)	(0.016)	(0.026)
Over target × Experienced	2104.639***	0.021	-0.010	-0.007
	(654.821)	(0.030)	(0.020)	(0.033)
Female CV	279.118	0.028	-0.027*	0.015
	(331.832)	(0.023)	(0.014)	(0.022)
Female CV × Experienced	-761.177	0.017	0.014	0.022
	(489.876)	(0.032)	(0.018)	(0.032)
Over target × Female CV	-946.346	-0.042	0.032	0.046
	(681.131)	(0.033)	(0.021)	(0.031)
Over target × Female CV × Experienced	1007.449	-0.037	-0.007	-0.056
	(995.662)	(0.043)	(0.027)	(0.044)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749

Notes: This table examines how recruiters evaluate candidates depending on their graduation year (Panel A) and their experience level (Panel B), with differential effects by gender. Candidates' age is proxied by their year of graduation: young candidates graduated between 2016-and 2022, versus 2007-2015 for the older candidates. Experienced candidates have above-median average years of experience (4.5 years or more). Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B12 – Effects of Male-Favoring Recruiters

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
Male favoring recruiters	934.409*	0.064	-0.026	-0.011
	(523.579)	(0.040)	(0.024)	(0.045)
Over target	14723.955***	0.163***	0.050*	0.012
	(833.789)	(0.036)	(0.027)	(0.039)
Over target × Male favoring recruiters	-327.869	-0.055	-0.031	0.046
	(1014.168)	(0.044)	(0.032)	(0.047)
Female CV	520.167	0.072**	-0.035*	0.021
	(456.280)	(0.035)	(0.021)	(0.030)
Female CV × Male favoring recruiters	-872.739	-0.053	0.007	-0.002
	(546.861)	(0.041)	(0.024)	(0.035)
Over target × Female CV	556.642	-0.024	-0.001	0.042
	(858.689)	(0.047)	(0.032)	(0.044)
Over target × Female CV × Male favoring recruiters	-2154.368**	-0.055	0.065*	-0.033
	(1048.613)	(0.056)	(0.037)	(0.053)
Y mean	144,991	.57	.12	.59
N	5,927	5,927	5,927	5,927

Notes: This table examines how recruiters who favor male candidates evaluate candidates differently. This figure displays the coefficients for whether the recruiter is designated as male-favoring. See notes in Figure B3 for constructing the variable for male-favoring propensity. Male-favoring recruiters are defined as those who have the propensity greater the median. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B13 – Effects of Recruiters with High MSA-Level Gender Bias

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
Panel A: Gender – Career				
Gender Bias	-202.610	-0.019	0.018	0.022
	(515.765)	(0.038)	(0.020)	(0.046)
Over target	14633.368***	0.068**	0.058***	0.038
	(694.336)	(0.028)	(0.018)	(0.032)
Over target × Gender Bias	717.998	0.085**	-0.031	0.038
	(977.611)	(0.040)	(0.026)	(0.045)
Female CV	-360.098	-0.015	0.001	0.025
	(397.354)	(0.026)	(0.014)	(0.021)
Female CV × Gender Bias	506.497	0.084**	-0.034*	-0.012
	(537.544)	(0.036)	(0.020)	(0.031)
Over target × Female CV	113.036	0.010	-0.006	0.047
	(742.707)	(0.038)	(0.022)	(0.033)
Over target × Female CV × Gender Bias	-1715.177*	-0.122**	0.047	-0.028
	(1035.121)	(0.054)	(0.033)	(0.049)
Y mean	145,728	.6	.1	.63
N	5,571	5,571	5,571	5,571
Panel B: Gender – Science				
Gender Bias	-180.134	-0.003	-0.001	-0.026
	(509.949)	(0.038)	(0.019)	(0.045)
Over target	14476.233***	0.068**	0.055***	0.038
	(628.425)	(0.027)	(0.018)	(0.032)
Over target × Gender Bias	1059.233	0.089**	-0.026	0.036
	(978.146)	(0.040)	(0.026)	(0.045)
Female CV	-341.084	-0.012	0.001	-0.023
	(370.892)	(0.026)	(0.015)	(0.022)
Female CV × Gender Bias	466.992	0.077**	-0.033	0.082***
	(532.305)	(0.036)	(0.020)	(0.031)
Over target × Female CV	-212.923	-0.010	0.005	0.082**
	(696.342)	(0.037)	(0.023)	(0.034)
Over target × Female CV × Gender Bias	-1061.865	-0.081	0.025	-0.094*
	(1036.230)	(0.054)	(0.033)	(0.049)
Y mean	145,728	.6	.1	.63
N	5,571	5,571	5,571	5,571

Notes: This table examines whether candidates face differential negotiation penalties depending on the gender bias prevalent in the recruiter's MSA. Coefficients reported are for the interaction term indicating a high-bias MSA. MSAs are classified as high-bias if their average Gender-Science IAT score exceeds the sample median. The Gender–Science IAT captures implicit associations between gender and fields (e.g., science versus liberal arts), with higher values indicating stronger stereotypical associations linking males with science. IAT scores are averaged at the MSA level using respondents from 2015–2020. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B14 – Effects of Female Recruiters

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
Female recruiters	302.438	0.010	-0.002	0.117***
	(483.931)	(0.038)	(0.020)	(0.041)
Over target	15992.821***	0.132***	0.047**	0.114***
	(843.628)	(0.035)	(0.023)	(0.036)
Over target × Female recruiters	-1612.430*	-0.015	-0.023	-0.089**
	(974.027)	(0.040)	(0.026)	(0.042)
Female CV	35.818	0.054*	-0.018	0.070***
	(475.710)	(0.031)	(0.019)	(0.025)
Female CV × Female recruiters	-224.809	-0.024	-0.004	-0.059**
	(534.152)	(0.036)	(0.022)	(0.029)
Over target × Female CV	396.806	-0.077*	0.022	-0.036
	(861.602)	(0.045)	(0.027)	(0.042)
Over target × Female CV × Female recruiters	-1184.367	0.019	0.009	0.072
	(1004.005)	(0.052)	(0.032)	(0.049)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749

Notes: This table examines how male and female recruiters differentially perceive and evaluate candidates of different gender. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B15 – Spillover Effects of Other Negotiating Candidates

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
High % over target	820.806**	-0.019	-0.039**	-0.008
	(414.903)	(0.032)	(0.017)	(0.038)
Over target	12570.384***	0.104***	0.033	0.059**
	(892.105)	(0.028)	(0.023)	(0.028)
Over target × High % over target	2338.925**	0.028	0.011	-0.008
	(1013.120)	(0.035)	(0.026)	(0.035)
Female CV	83.598	0.039*	-0.032**	0.038**
	(293.333)	(0.021)	(0.013)	(0.017)
Female CV × High % over target	-411.858	-0.006	0.024	-0.023
	(436.512)	(0.031)	(0.017)	(0.027)
Over target × Female CV	857.698	-0.068	0.028	0.012
	(1116.616)	(0.044)	(0.031)	(0.039)
Over target × Female CV × High % over target	-1415.112	0.008	-0.009	0.015
	(1236.835)	(0.052)	(0.035)	(0.049)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749

Notes: This table examines how recruiters evaluate candidates varying the share of candidates who actively negotiate, with differential effects by gender. In the high-share group, more than 50% candidates or more negotiate over the salary target. In the low-share group, 25% of candidates negotiate over the salary target. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

Table B16 – Spillover Effects of Other Negotiating Female Candidates

	Salary Offer	Recommend for Senior Position	Reject Application	Work from Home
High Share Women Negotiate	598.007	-0.035	-0.018	-0.024
	(441.526)	(0.033)	(0.016)	(0.039)
Over target	14176.390***	0.119***	0.023	0.056**
	(580.544)	(0.022)	(0.015)	(0.024)
Over target × High Share Women Negotiate	865.119	0.019	0.020	-0.001
	(845.524)	(0.034)	(0.021)	(0.037)
Female CV	-3.349	0.021	-0.024**	0.018
	(244.851)	(0.018)	(0.010)	(0.015)
Female CV × High Share Women Negotiate	-344.221	0.055	0.009	0.031
	(527.299)	(0.035)	(0.020)	(0.031)
Over target × Female CV	-467.809	-0.065**	0.047**	0.005
	(667.368)	(0.031)	(0.019)	(0.027)
Over target × Female CV × High Share Women Negotiate	-21.782	-0.025	-0.033	0.005
	(940.657)	(0.048)	(0.029)	(0.046)
Y mean	145,638	.62	.099	.62
N	7,749	7,749	7,749	7,749

Notes: This table examines how recruiters evaluate candidates varying the share of female candidates who actively negotiate, with differential effects by gender. In the high-share group, more than 25% female candidates negotiate over the salary target. In the low-share group, 25% of female candidates or less negotiate over the salary target. Except for salary estimates, all variables are binary, coded as 1 if they exceed the midpoint on a 4-point scale. All regressions include biography fixed effects and recruiter controls, with standard errors clustered at the recruiter level.

C Structural Estimation

We propose a structural estimation of the model presented in Section 1. A single set of firm-side primitives—the recruiter's beliefs about the candidate, the value of skill and seniority, rigidities in pay and seniority, and a residual term—generates all of these responses at once, and we estimate them jointly by matching six survey outcomes per cell (perceived skill, outside option, wage, seniority, acceptance, and rejection). Each response differs systematically by gender, and we put the estimates to two uses. Section C.3 decomposes the gender gaps into a *belief* channel—statistical discrimination in how skill is perceived—and a *residual* channel—a gender- and ask-specific term in the firm's objective, whose source we do not take a stand on. Section C.4 then inverts the estimated firm to derive the candidate's optimal counteroffer. Finally, we compare how the estimated model parameters shift under pay secrecy versus transparency.

C.1 Estimation

C.1.1 Revised Model

A candidate in cell c is summarized by recruiter beliefs $(\hat{v}_c, \hat{\omega}_c)$: perceived skill and perceived outside option. The recruiter (firm) chooses a wage w and a seniority recommendation $s \in [0,1]$ to maximize expected profit

$$\Pi(w,s;c) = \mathbf{\Lambda}_{\text{struct}}\left(\tfrac{w+\tau s-\hat{\omega}_c}{\sigma_{\text{struct}}}\right) B(w,s;c) - c_R,$$

where $\mathbf{\Lambda}_{\text{struct}}$ is the probability the candidate accepts and the surplus if accepted is redefined as:

$$B = \pi_v \hat{v}_c(1+\delta s) - w - \tfrac{1}{2}\kappa_w(w-\bar{w}_{\text{eff}})^2 - \tfrac{1}{2}\kappa_s\, s^2 - g_s\, s - g_w(w-\bar{w}_{\text{eff}}) + c_R,$$

Residual shifters. Where the rigidities κ_w, κ_s are gender- and ask-neutral—governing the firm's response to its *beliefs* about a candidate, identically for everyone—the shifters g_s, g_w add a *systematic* gender $\times$ ask tilt in the marginal value of seniority and wage, holding beliefs fixed (they enter B as $-g_s s$ and $-g_w(w-\bar{w}_{\text{eff}})$). Writing $i \in \{s,w\}$, both share one saturated

form,

$$g_i = \tfrac{1}{15}\Big[\gamma^i_{fo}\,F(a{-}155)_+ + \gamma^i_{fu}\,F(155{-}a)_+ + \gamma^i_{mo}\,M(a{-}155)_+ + \gamma^i_{mu}\,M(155{-}a)_+\Big],$$

where $F, M \in \{0,1\}$ are gender indicators, a the ask, 155 (\$K) the posted-range top (the social-norm reference $a^\star$), and $(\cdot)_+$ the positive part. The divisor 15 is a units normalization—each γ is the effect, in \$K, of a \$15K over- or under-ask—and rescales the coefficients without affecting fit or inference. This gives eight coefficients, four per margin; we drop the superscript on the focal seniority margin ($\gamma \equiv \gamma^s$) and keep γ^w for the wage, estimated jointly.

Acceptance and the extension margin. The firm's problem is two-stage. Given the candidate, it first chooses $(w^\star, s^\star)$ to maximize the expected profit Π above; the candidate accepts with probability $\Lambda_{\text{struct}}\Big((w + \tau s - \hat{\omega}_c)/\sigma_{\text{struct}}\Big)$, rising in the total package $w + \tau s$ relative to the perceived outside option $\hat{\omega}_c$, with σ_{struct} the acceptance-noise scale. The maximized profit is $\Pi^\star = \Lambda_{\text{struct}}\,B(w^\star, s^\star) - c_R$: the acceptance probability times the surplus, net of the cost c_R of a declined offer. The firm then decides whether to extend an offer at all, and does so *probabilistically*,

$$\Pr(\text{extend}) = \Lambda\Big(\Pi^\star/\sigma_R\Big), \qquad \Lambda(z) = \tfrac{1}{1+e^{-z}},$$

a random-utility extensive margin: the firm extends when the maximized profit plus an idiosyncratic shock (competing candidates, budget, fit; logistic with scale σ_R) is positive. A profitable candidate ($\Pi^\star \gg 0$) is almost surely offered, an unprofitable one almost surely not, and σ_R governs how sharply the decision tracks profit. The candidate is hired only if the firm extends *and* the candidate accepts, so $P_{\text{hire}} = \Pr(\text{extend}) \cdot \Lambda_{\text{struct}}$, and the model's reject outcome is $1 - \Pr(\text{extend})$.

C.1.2 Estimation and Inference

Data and cells. The experiment randomizes two dimensions seen by the recruiter: the candidate's gender and the size of the candidate's counteroffer. For each evaluation the recruiter reports six outcomes—the candidate's perceived skill and outside option, a wage offer, a seniority recommendation, how likely the candidate is to accept, and whether they would decline to make an offer. We summarize the data by binning evaluations into ten cells: the two ran-

domized genders × five within-gender quantiles of the ask. Gender is the treatment of interest; the counteroffer is a continuous randomized margin, and quintiles let each outcome trace the counteroffer distribution—and the \$155K over-/under-target threshold—flexibly, without imposing a parametric ask gradient. Five quantiles trade resolution (enough to capture the belief and response gradients and the over-versus-under split) against cell size; with $N = 7{,}733$ evaluations by 492 recruiters, each cell averages several hundred observations.

Variable construction and calibration. Perceived skill is the recruiter's composite endorsement across three dimensions—technical, teamwork, and initiative. The seniority outcome is an indicator that the recruiter recommends the candidate for the senior role, collapsing the four-level recommendation item to "strongly or moderately recommend" (sample mean 0.62). Wages, the outside option, and the counteroffer are the recruiter's dollar figures (in \$K). Three quantities are calibrated rather than estimated: the social-norm reference ask $a^\star = \$155\text{K}$ is the top of the posted salary range; the institutional wage anchor $\bar{w}_{\text{inst}} = \$145.4\text{K}$ is set to the sample-mean wage offer (Section C.1.1); and the survey-acceptance measurement scale is normalized, $\sigma_{\text{acc}} = 1$. All remaining parameters are estimated.

Moments and estimator. The ten cells × six outcomes give 60 unconditional moments (cell means), matched by 38 parameters: 20 are the cell beliefs $(\hat{v}_c, \hat{\omega}_c)$ themselves, which absorb the perceived-skill and outside-option moments. The belief block is therefore nonparametric—the $\hat{v}_c, \hat{\omega}_c$ are the fitted values of a saturated regression of the two survey measures on gender × ask-quintile, so the model places no structure on beliefs, only on the firm's response to them. The 18 firm-side parameters (ten firm-side plus eight γ) must then fit the remaining 40 moments—wage, seniority, acceptance, and rejection across the ten cells—leaving 22 over-identifying restrictions. Estimation is classical minimum distance with a diagonal inverse-variance weight,

$$Q(\theta) = \sum_{c=1}^{10} \sum_{k=1}^{6} \frac{(\bar{m}_{c,k}^{\text{data}} - m_{c,k}^{\text{model}}(\theta))^2}{\widehat{\text{Var}}(\bar{m}_{c,k}^{\text{data}})},$$

minimized by L-BFGS-B with multistart. Standard errors are a recruiter-clustered bootstrap (cluster on 492 recruiters, $B = 200$, re-binned and re-estimated on each draw).

C.2 Parameter Estimates

Table C1 reports the full-sample estimates. Production is skill-driven and complementary with seniority: $\pi_v \approx 217$ scales the perceived-skill index into output, and $\delta = 0.27$ ($[0.26, 0.29]$) means a top-seniority placement makes perceived skill 27% more valuable—the cross-partial through which a biased belief about skill becomes a seniority gap. Rigidities are sticky and asymmetric: the seniority cost is large ($\kappa_s \approx 88$), and the wage rigidity ($\kappa_w \approx 11$)[A.1] holds offers near an anchor set almost entirely by the candidate's own market value ($\alpha = 0.91$, so $dw^\star/d\hat{\omega} \approx \alpha$). The residual term is estimated jointly with γ^w for the wage margin, since both feed into the shared acceptance and extension moments; all eight coefficients are individually significant, which we use below (Section C.3) to construct the belief-equalization counterfactual.

δ and κ_s are identified jointly off the seniority margin—κ_s from the level of $s^\star$, δ from its cross-cell slope in perceived skill ($\kappa_s s^\star \approx \pi_v \delta \hat{v} - g_s$). With the cross-cell spread in $\hat{v}$ small (SD ≈ 0.02) and no exogenous skill shifter in the experiment, this slope is weakly identified and δ's magnitude leans on the $(1 + \delta s)$ form; we read δ as a transmission coefficient, not a sharp primitive. However, none of our conclusions turns on its magnitude: the belief asymmetry $\Delta^2 \hat{v}$ (the gender $\times$ over-vs-low-counteroffer difference-in-differences in perceived skill) is model-free and identified off randomized gender, the belief channel of the decomposition is significant whatever δ, and the pay-transparency result is identified off randomized gender $\times$ ask with δ flat across regimes.

Figure C1 shows clear patterns of the effects of more intense negotiation on skill and outside option evaluations. Male perceived skill $\hat{v}$ rises with the counteroffer while female $\hat{v}$ plateaus—the asymmetry $\Delta^2 \hat{v} = 0.041$. The outside option $\hat{\omega}$, by contrast, rises with the ask for both genders in close parallel. The differential treatment is specific to skill evaluation: recruiters credit men's high counteroffers with more skill, not with better outside options.

Fit. Table C2 compares data and model on each outcome's cell mean and gender DiD, and Figure C2 plots the full cell-by-cell fit. The model reproduces the levels almost exactly (mean

[A.1] The raw values of κ_w and κ_s are not directly comparable, since w and s live on very different scales (wages vary by tens of thousands of dollars; s is bounded on $[0,1]$). Measured per standard deviation of each variable—the scale-invariant comparison underlying Appendix A.4—wage-setting is in fact far more rigid than senior-track placement: $\kappa_w \, \mathrm{sd}(w)^2 \approx 1{,}870$ versus $\kappa_s \, \mathrm{sd}(s)^2 \approx 21$, consistent with the condition favoring stronger senior responsiveness to skill.

seniority 0.619 vs 0.619, mean wage \$145.8K vs \$145.8K, mean rejection 0.099 vs 0.098), which is unsurprising for the two belief outcomes—the cell beliefs absorb them—but is a genuine restriction for wage, seniority, and the two probabilities, which the eighteen firm-side parameters must fit jointly. More demanding are the gender × counteroffer difference-in-differences. The model matches the headline ones: the belief-asymmetry first stage (skill DiD +0.042 vs data +0.041), the senior gender gap (+0.063 vs +0.081, a gap of less than one bootstrap SE), and the rejection gap (−0.044 vs −0.037). The wage DiD is the one place the model overshoots (+0.76 vs +0.57), but with a bootstrap SE of about 0.45 the discrepancy is well inside one standard error—this wage DiD is imprecisely estimated in the data (the data value itself carries an SE of 0.43)— so it is not a meaningful failure and does not bear on the seniority results. Overall the model is consistent with every targeted moment within sampling error.

C.3 Decomposing The Gender Gap

Having estimated the model and shown it fits, we use it to ask what drives the gender gaps it reproduces—how much each reflects biased beliefs versus residual, and how these mechanisms explain the effect of pay transparency.

We are interested in three firm-side outcomes that show a gender gap—the wage offer, the seniority recommendation, and the rejection rate—and we decompose each with a common method. For an outcome y, define its gender difference-in-differences, the differential response to over versus below target counteroffer, men relative to women,

$$\mathrm{DiD}(y) = \Big[\bar{y}_{M,\text{over}} - \bar{y}_{M,\text{under}}\Big] - \Big[\bar{y}_{F,\text{over}} - \bar{y}_{F,\text{under}}\Big],$$

where "over"/"under" is relative to the \$155K posted-range top and cell means are count-weighted. We split each DiD into two channels by a nested counterfactual. First we equalize beliefs across genders—replacing each counteroffer quintile's $(\hat{v},\hat{\omega})$ with its gender-pooled average, so a man and a woman in the same quintile face identical perceived skill and outside option—and recompute the DiD from the firm's first-order conditions. With beliefs now gender-symmetric, the only thing left that can produce a gender DiD is the gender-asymmetric residual penalty. The *belief-*×*-*δ channel is the remainder—i.e. the part of the gap the gender

belief asymmetry generates, operating through the firm's belief-driven response (production $\pi_v \hat{v}$, the wage anchor $\alpha\hat{\omega}$, and—for seniority—the complementarity δ).

Table C4 reports the decomposition for three outcomes—the wage offer, the seniority recommendation, and the rejection rate. The pattern is consistent: the belief-$\times$-δ channel is individually significant for seniority and rejection, and marginally so for the wage (its CI just reaches zero), while the residual channel (γ) is insignificant for all three in the full sample. The belief asymmetry is thus a pervasive first stage—the same biased skill belief pushes high-counteroffer men's wages, seniority, and hiring up. For seniority decision, only the belief-$\times$-δ channel is individually significant ($+0.040$, CI $[+0.013, +0.064]$); the residual channel is positive but insignificant ($+0.022$, CI $[-0.012, +0.076]$). Because these two CIs overlap and δ trades off with γ on the firm side, we don't read a sharp share split—the point estimates put roughly two-thirds on belief-$\times$-δ, but that ratio is imprecise. What the data support is the sign-and-significance pattern: the senior gap is carried by the channel tied to biased perceived skill, with the residual channel not separately identified in the pooled sample—it emerges sharply, and for seniority specifically, once we split by pay transparency.

C.4 Optimal Counteroffers

The estimated model also rationalizes a second gender gap: the gap in how aggressively candidates should counteroffer. We invert the estimated firm to ask what counteroffer a candidate would optimally make. Treating the firm's response as the candidate's choice environment, a candidate of gender g chooses the ask c to maximize expected utility

$$\mathrm{EU}(c) = P_{\text{hire}}(c, g)\, u\Big(w^\star(c, g) + \tau s^\star(c, g)\Big),$$

where $P_{\text{hire}} = P_{\text{extend}} \cdot \Lambda_{\text{struct}}$ is the hiring probability the estimated firm implies at counteroffer c, and $w^\star + \tau s^\star$ is the resulting compensation if hired. If the counteroffer is rejected the candidate loses the offer, so the not-hired payoff is zero: over-asking trades a richer package against a lower chance of being hired, which gives an interior optimum. We compute $c^\star = \arg\max_{c \in [\$120, \$200]\text{K}} \mathrm{EU}(c)$ under three preference specifications. The first is risk neutrality, $u(x) = x$. The second is constant relative risk aversion, $u(x) = (x^{1-\rho} - 1)/(1 - \rho)$, evaluated at $\rho = 0.5$. The third is the reference-dependent value function of Tversky and Kahneman

(1992),

$$v(x) = \begin{cases} (x-r)^{\alpha}, & x \geq r, \\ -\lambda\,(r-x)^{\alpha}, & x < r, \end{cases}$$

with reference point $r = \$130\text{K}$ (the initial bid) and their estimated parameters $\alpha = 0.88$ and $\lambda = 2.25$.

We find that the optimal counteroffer is higher for men, and the gap is belief-driven. Figure 3 and corresponding Table C5 report the optimal counteroffer by gender. Men optimally renegotiate higher than women under every utility specification. Women's optimal counteroffer is stable around \$146 (about +12% over the \$130K initial bid) regardless of risk preference; men's is \$157 (about +21%) under curvature and \$168 (about +30%) under risk neutrality, so the gap compresses from +17pp under risk neutrality to about +9pp under curvature. Figure C3 shows the underlying sweep: reject and hiring probabilities, compensation, and expected utility by gender.

These gaps are driven by differences in hiring probability, which traces directly to the belief asymmetry of Section C.2. Figure C3 shows the mechanism directly: the hiring probability falls with the counteroffer for both genders but far more steeply for women, while compensation conditional on hire rises smoothly and similarly for both. Table C6 gives the anatomy behind the figure, walking the firm's response from a low counteroffer (quintile 1) to a high counteroffer (quintile 5) for each gender. As the counteroffer rises, male perceived skill climbs steeply ($\Delta\hat{v} = +0.036$) while female perceived skill is essentially flat ($\Delta\hat{v} = +0.002$), the same $\Delta^2\hat{v}$ asymmetry that drives the seniority gap. Because a higher-asking man is perceived as more skilled, his surplus to the firm erodes far less as he renegotiates more: the male hiring probability falls only 5.5 percentage points across the counteroffer range, versus 8.1 for women, and the male surplus B falls \$6.9K versus \$16.3K.

Splitting by regime (Table **??**) shows that the optimal-counteroffer gap echoes the seniority gap: it is wide under pay secrecy, and compresses toward its belief-driven core under transparency. Under transparency the male optimum is interior in every utility specification, and the gap—+15.4pp under risk neutrality and +8.5pp under curvature—closely tracks the full-sample estimates (+17.3 and +8.8pp); equalizing beliefs sends it to zero, exactly as in the pooled sample. Under secrecy the gap is wider, and its risk-neutral magnitude is not sharply

defined: the male-under role-conformity penalty that transparency removes (Table C8) reinforces the complementarity δ in pushing men's counteroffer upward, and for a risk-neutral candidate this censors the optimum at the \$200K grid ceiling, so the reported +41.5pp is a lower bound on the true gap. Where the optimum remains interior—under both curvature specifications—the secrecy gap is a more moderate +7.7pp.

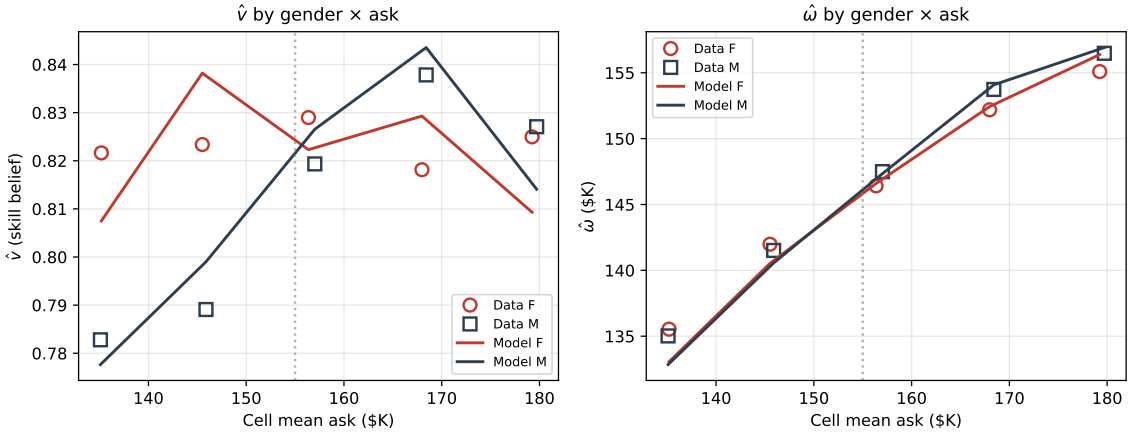
The regime split therefore reinforces the seniority-gap decomposition: transparency narrows the optimal-counteroffer gap through the same channel by which it narrows the seniority gap. In every specification, men's optimal counteroffer exceeds women's, and in every specification transparency moves the gap toward the magnitude implied by beliefs alone.

C.5 Pay transparency

We re-estimate the model separately under pay secrecy and transparency, randomly assigned across evaluations, and compare the two parameter vectors (Table C8). The seniority cost κ_s rises under transparency, a suggestive pattern ($p = 0.11$): the point estimate is consistent with the firm becoming more broadly cautious about senior-track recommendations under transparency, for men and women alike. The remaining parameters, including the complementarity δ, the wage anchor α, and the belief-driven core, are statistically unchanged across regimes.

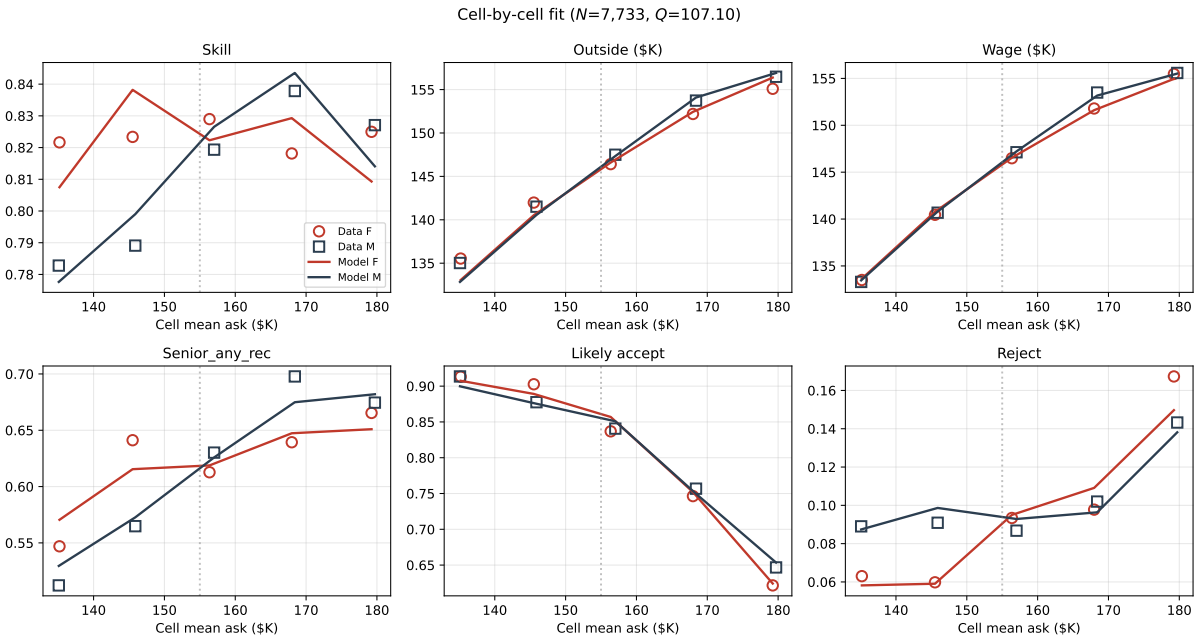
The same pattern shows up when we decompose the senior gap itself within each regime (Table C9). The senior gender DiD is large and significant under secrecy (+0.097, $p < 0.01$) but compresses to a statistically insignificant +0.035 under transparency, even as the belief channel remains essentially unchanged across regimes (+0.037 under secrecy versus +0.043 under transparency). The compression in the gap is therefore not driven by any change in how recruiters' beliefs translate into seniority recommendations; it coincides instead with the (suggestive) broader tightening in κ_s, consistent with transparency operating through firm discretion generally rather than through a channel specific to gender. It is also interesting to note that under secrecy the unexplained residual channel (γ) carries the majority of the gap (+0.061, 61%)

FIGURE C1. Perceived Skill $\hat{v}$ and Outside Option $\hat{\omega}$ by Gender and Counteroffer



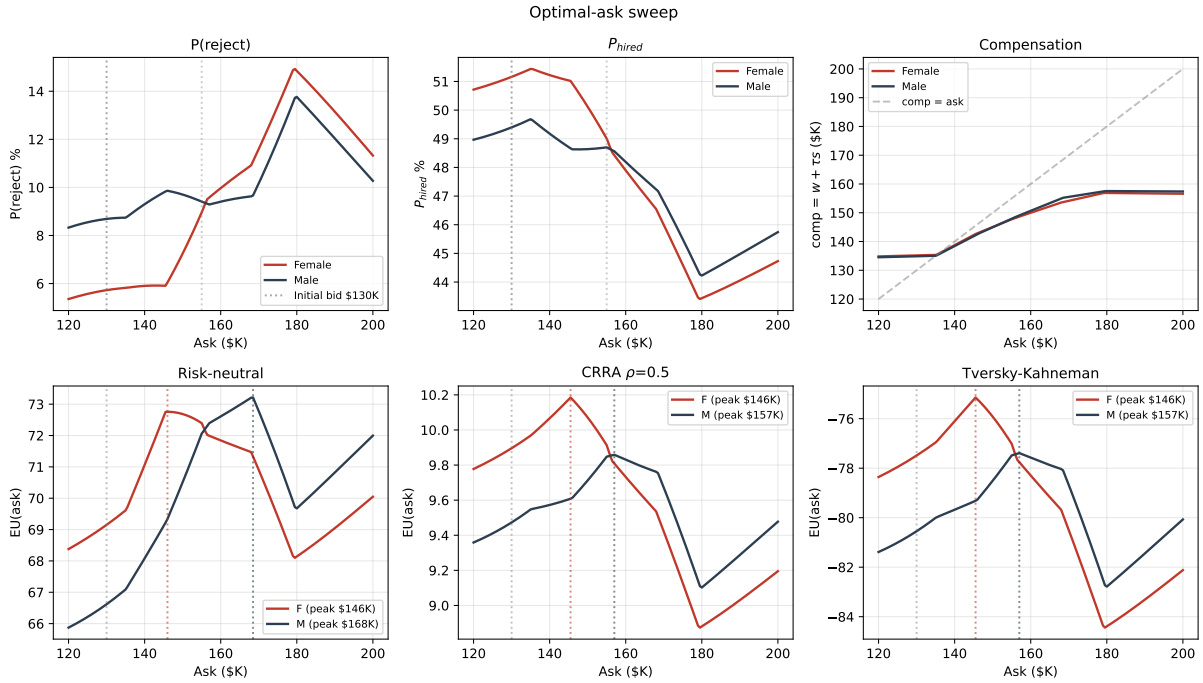
Notes: The figure shows the model fit, by plotting the evolution of the perceived skill $\hat{v}$ and outside option $\hat{\omega}$ by gender and counteroffer distribution. Markers are data cell means; lines connect the model's cell estimates (five ask quantiles per gender). Estimation on the full sample ($N = 7{,}733$, $Q = 107.10$).

FIGURE C2. Cell-by-Cell Fit



Notes: The figure present cell-by-cell model fit across the six outcomes of interest of the recruiter, namely the evaluation of skill, outside option, recommended wage, recommendation for senior-track position, perceived likelihood that the candidate will accept the offer, and rejection of the candidate's application. Markers are data cell means; lines connect the model's cell estimates (five ask quantiles per gender). Estimation on the full sample ($N = 7{,}733$, $Q = 107.10$).

FIGURE C3. Optimal Counteroffers



Notes: Each panel plots a quantity against the candidate's counteroffer $c \in [\$120, \$200]$K, separately for men and women, from the candidate's expected utility against the estimated firm from Equation (11). *P(reject)* is the chance the recruiter recommends rejecting the candidacy; P_{hire}= $P_{\text{ext}} \times$ (acceptance probability) is the chance the candidate is both extended an offer and accepts it (so $P_{\text{hire}} \leq 1 - P(\text{reject}) = P_{\text{ext}}$, the wedge being the acceptance margin); compensation is $w^\star + \tau s^\star$ conditional on hire; and expected utility is $EU(c) = P_{\text{hire}} \cdot u(compensation)$, shown per utility specification. The counteroffer that maximizes expected utility is the optimal counteroffer $c^\star$, marked by the peak of the expected-utility panel.

Table C1 – Structural parameter estimates

Parameter	Estimate	95% CI
Production and rigidity		
δ	0.2737	[+0.2592, +0.2864]
π_v	216.78	[+214.11, +217.09]
κ_w	10.84	[+9.43, +11.93]
κ_s	88.19	[+85.50, +88.44]
α	0.91	[+0.89, +0.92]
τ	2.98	[+2.77, +3.37]
c_R	4.01	[+3.17, +4.80]
Noise (acceptance / extension)		
σ_{struct}	12.30	[+12.00, +13.46]
σ_R	9.92	[+8.84, +10.81]
σ_{acc}	1.00	[+1.00, +1.00]
Residual shifters (senior)		
γ_{fo} (F-Over)	−2.75	[−4.27, −0.53]
γ_{fu} (F-Under)	+3.01	[+1.65, +5.51]
γ_{mo} (M-Over)	−4.09	[−5.62, −2.48]
γ_{mu} (M-Under)	+3.75	[+2.66, +6.79]
Residual shifters (wage)		
γ^w_{fo}	+3.72	[+2.27, +4.90]
γ^w_{fu}	+4.74	[+3.73, +6.92]
γ^w_{mo}	+3.16	[+1.60, +4.56]
γ^w_{mu}	+4.39	[+3.12, +6.55]
N	7,733	
Q at $\hat{\theta}$	107.10	
Bootstrap draws	200/200 converged	

Notes: This table shows all the parameters of the structural model described in Section 6. Details of the procedure are presented in Appendix Section C. The estimation sample is the full experimental sample, $N = 7{,}733$. 95% confidence-intervals are obtained from a recruiter-clustered bootstrap ($B = 200$). σ_{acc} normalized.

Table C2 – Model Fit: Data vs. Model

	Cell mean		Gender DiD	
Outcome	Data	Model	Data	Model
Perceived skill	0.817	0.816	+0.041 (0.017)	+0.041 (0.014)
Outside option ($K)	146.4	145.9	+1.62 (0.69)	+0.81 (0.47)
Wage offer ($K)	145.6	145.6	+0.57 (0.43)	+0.76 (0.45)
Seniority recommendation	0.617	0.618	+0.081 (0.024)	+0.062 (0.022)
Acceptance prob.	0.807	0.807	+0.025 (0.017)	+0.017 (0.016)
Rejection prob.	0.099	0.098	−0.037 (0.014)	−0.044 (0.014)

Notes: This table presents average moments in the data (full sample) and in the model for the outcomes of interest. The model moments are computed for estimated values of the structural parameters. Cell means and the gender × counteroffer difference-in-differences (over- minus under-target, men minus women, count-weighted). Recruiter-clustered bootstrap standard errors ($B = 200$) on the DiDs are in parentheses, computed on the same resamples for data and model. The two belief outcomes (skill, outside option) are absorbed by the cell-belief parameters; wage, seniority, acceptance, and rejection are fit by the eighteen recruiter-side parameters.

Table C3 – Model Misfits: Impact on the Decomposition

Panel A. Contribution to the criterion Q, by outcome

Outcome	Share of Q	Largest $\lvert t\rvert$
Perceived outside option	65.4%	4.03
Perceived skill	8.6%	1.13
Seniority recommendation	7.8%	1.26
Wage offer	7.2%	1.33
Acceptance probability	6.7%	1.39
Rejection probability	4.3%	1.15

Panel B. Sensitivity of the decomposition channels (share of identification)

Channel	Outside option	Skill	Seniority
Belief $\times\,\delta$	3.1%	27.3%	24.1%
Residual (γ)	3.7%	22.7%	40.3%

Notes: Panel A gives each outcome's contribution (summed over the ten gender×ask cells) to the minimum-distance criterion at $\hat{\theta}$; $Q(\hat{\theta}) = 107.10$ under the diagonal inverse-variance weight used in estimation, and $Q = 68.47$ under a recruiter-cluster-robust reweighting. "Largest $|t|$" is the largest across cells of the cluster-robust standardized residual (data − model)/SE. The misfit concentrates in the perceived outside option in the two lowest-ask cells (F_{q1}, M_{q1}: model \$133.0K vs. data \$135.5K, a 1.8% miss); all other outcomes fit within ≈ 1.4 SE. Panel B reports the sensitivity of the belief and the residuals to the empirical moments (Andrews, Gentzkow and Shapiro, 2017), as each moment group's share of the channel's standardized |sensitivity|; the outcome carrying 65% of Q supplies $\approx 3\%$ of each channel's identification, so the misfit is orthogonal to the decomposition.

Table C4 – Channel Decomposition by Outcome (Full Sample)

Outcome	Channel	Contribution	95% CI
Wage offer ($K)	total	+0.76	[−0.06, +1.68]
	belief×δ	+1.02	[+0.00, +1.84]
	γ	−0.26	[−0.41, +0.35]
Seniority	total	+0.062	[+0.026, +0.114]*
	belief×δ	+0.040	[+0.013, +0.064]*
	γ	+0.022	[−0.012, +0.076]
Rejection	total	−0.044	[−0.070, −0.017]*
	belief×δ	−0.037	[−0.065, −0.010]*
	γ	−0.006	[−0.023, +0.004]

Notes: This table presents the contribution of each channel to the gender × counteroffer difference-in-differences, by outcome, recomputed on every one of $B = 200$ recruiter-clustered draws. The row belief×δ corresponds to the total gap minus the γ channel; the γ channel corresponds to the difference-in-differences after equalizing beliefs. The star * indicates that the 95%-confidence intervals exclude zero.

Table C5 – Optimal Counteroffer by Utility Specification (Full Sample)

Specification	F^* (\$K)	M^* (\$K)	F^* % over	M^* % over	Gap (pp)
Risk-neutral	146.0	168.5	+12.3%	+29.6%	+17.3
CRRA ρ=0.5	145.5	157.0	+11.9%	+20.8%	+8.8
Tversky–Kahneman	145.5	157.0	+11.9%	+20.8%	+8.8

Notes: This table presents the optimal counteroffer $c^\star$ maximizing the candidate's expected utility against the estimated firm from Equation (11), by gender and risk preference (full sample), across utility specifications: risk-neutrality (first row), CRRA ρ=0.5, and a reference-dependent value function of Tversky and Kahneman (1992) with reference point r=\$130K and estimated parameters $\alpha = 0.88$ and $\lambda = 2.25$. The column "% over" is relative to the \$130K initial bid. The column "gap" is the gender difference in optimal counteroffer $M^\star - F^\star$ in percentage points. †: optimum at a bound of the \$120–200K ask grid.

Table C6 – Optimal Counteroffers: Mechanisms

Cell	$\hat{v}$	$\hat{\omega}$	$s^\star$	P_{ext}	P_{hire}	B (\$K)
Female, Q1 (low ask)	0.807	133.1	0.570	0.942	0.514	57.9
Female, Q5 (high ask)	0.809	156.4	0.651	0.850	0.434	41.6
Δ Female (pp for prob.)	+0.002	+23	+0.080	−9.2	−8.1	−16.3
Male, Q1 (low ask)	0.778	132.8	0.530	0.913	0.497	50.1
Male, Q5 (high ask)	0.814	156.9	0.682	0.862	0.442	43.2
Δ Male (pp for prob.)	+0.036	+24	+0.152	−5.1	−5.5	−6.9

Notes: This table examines the mechanisms behind the differences in optimal counteroffers by gender. Each gender has three rows: the estimated recruiter's response when the candidate renegotiates at the lowest (Q1) and highest (Q5) counteroffer quintile, and the difference Δ = Q5−Q1. Evaluated at the cell ask, the columns are the recruiter's perceived skill $\hat{v}$ and perceived outside option $\hat{\omega}$ (\$K); the seniority the recruiter optimally recommends $s^\star$; the probability the recruiter extends an offer P_{ext}; the probability the candidate is hired $P_{\text{hire}} = P_{\text{ext}} \times$ (acceptance probability); and the firm's surplus B (\$K). P_{ext} is a firm-side decision, while P_{hire} also requires the candidate to accept the offer, so $P_{\text{hire}} \leq P_{\text{ext}}$ and the wedge between them is the acceptance margin. In the Δ rows, probability changes are in percentage points and $\hat{\omega}$ and B in \$K.

Table C7 – Optimal-Counteroffer Gender Gap (pp): Counterfactual Decomposition

Scenario	Risk-neutral	CRRA ρ=0.5	T–K
Baseline	+17.3	+8.8	+8.8
CF1: Equalize beliefs	+10.0	+0.4	+0.4
CF2: Zero penalties ($\gamma = 0$)	+17.3	−19.6	−19.6
CF3: Symmetrize penalties	+16.9	+8.8	+8.8
CF4: Both	+0.4	+0.0	+0.4

Notes: This table presents optimal-counteroffer gender gaps ($M^\star - F^\star$, in percentage points over the \$130K initial bid) under each counterfactual, computed from Equation (11), by gender and across utility specifications: risk-neutrality (first row), CRRA ρ=0.5, and a reference-dependent value function of Tversky and Kahneman (1992) with reference point r=\$130K and estimated parameters $\alpha = 0.88$ and $\lambda = 2.25$. CF1 pools beliefs across genders; CF2 zeroes all γ; CF3 makes the γ's gender-symmetric; CF4 combines CF1 and CF2.

Table C8 – Parameter Shifts Under Pay Transparency

Parameter	Secrecy (p0)	Transparency (p1)	Shift (p1−p0)	95% CI of shift	p
δ	0.276	0.275	−0.001	[−0.026, +0.020]	0.920
κ_w	10.40	10.83	+0.39	[−2.27, +2.86]	0.720
κ_s	87.76	88.28	+0.55	[−0.09, +3.45]	0.110
σ_{struct}	13.04	12.28	−0.62	[−2.07, +0.51]	0.340
α	0.912	0.911	−0.002	[−0.026, +0.020]	0.830
τ	3.11	2.96	−0.12	[−0.77, +0.47]	0.670
π_v	216.55	216.48	−0.08	[−0.99, +15.09]	0.870

Notes: This table presents the parameter shifts under pay secrecy or pay transparency. Selected parameters from the model are re-estimated separately under pay secrecy and transparency. "Shift" is the across-regime difference $\theta_{p1}-\theta_{p0}$; its 95%-confidence interval and p-value are from a paired recruiter-clustered bootstrap ($B = 200$, both regimes drawn on the same resample). ***/**/*: $p < 0.01/0.05/0.10$.

Table C9 – Senior-DiD Decomposition by Transparency Regime

Panel A. Pay Secrecy

Channel	Contribution	95% CI	Share [95% CI]
Baseline (model senior DiD)	$+0.0970^{***}$	$[+0.0477, +0.1650]$	—
Belief asymmetry $\times\delta$	$+0.0369^{*}$	$[-0.0015, +0.0743]$	$39\%\,[-2, 87]$
Residual (γ)	$+0.0612^{**}$	$[+0.0113, +0.1291]$	$61\%\,[13, 101]$

Panel B. Pay Transparency

Channel	Contribution	95% CI	Share [95% CI]
Baseline (model senior DiD)	$+0.0352$	$[-0.0327, +0.0876]$	—
Belief asymmetry $\times\delta$	$+0.0431^{**}$	$[+0.0057, +0.0768]$	$101\%\,[-460, 1376]$
Residual (γ)	-0.0086	$[-0.0627, +0.0381]$	$2\%\,[-1209, 554]$

Notes: This table presents the senior-gap decomposition of Table C4 computed within each pay-transparency regime, from the per-regime estimates. "Share" is each channel's percentage of the model senior difference-in-differences. Contributions, confidence intervals, and shares are computed from $B = 200$ recruiter-clustered draws. The stars $^{***}/^{**}/^{*}$ mark a contribution confidence interval excluding zero at the 1/5/10% level.

D Recruitment Materials

D.1 Job Posting for Recruiters

Help Set Compensation Packages for Candidates

We are a specialized software company dedicated to crafting AI-driven solutions. As a venture-backed, cutting-edge firm, we are experiencing rapid growth and have numerous exciting openings across multiple locations. Currently, we're in a mass hiring period for computer software engineer roles in Chicago, Boston, Austin, and Seattle. Given the scale of our hiring needs, we're seeking your expertise to finalize compensation packages for several key positions within our tech firm.

Responsibilities:

- Your participation in this task will likely require around 50 minutes of your time, followed by a brief 10-minute debriefing session.
- During this period, we ask you to review 10-20 candidate profiles and provide thoughtful evaluations using the evaluation form provided. Each evaluation will consist of around 10 short questions for each candidate. We expect you to offer insights and recommendations based on industry standards and candidate qualifications in response to our questions.

Qualifications:

- Expertise in HR, compensation setting, or a related field.
- Strong communication skills, both verbal and written.
- Ability to think critically about client needs and provide valuable advice from the client's perspective.
- Familiarity with current industry trends and best practices in compensation.
- Detail-oriented with a focus on accuracy and thoroughness.

If you're passionate about shaping the future of tech talent and have the expertise to help us navigate our hiring process, we'd love to hear from you. Join us in this exciting opportunity to contribute to our growth and success.

D.2 Recruitment Email

I hope this message finds you well. I'm reaching out with an opportunity for a short-term recruiting project that requires a brief 1-hour commitment with $50 compensation. My client, a specialized software company focused on AI-driven solutions, is currently in a mass hiring phase for computer software engineer roles across Chicago, Boston, Austin, and Seattle.

We're seeking someone with an HR/recruiting background to assist in finalizing compensation packages for key positions within the tech team. The task will take approximately 50 minutes, followed by a 10-minute debriefing session. This project is posted on [the platform name], where you would work through the platform. If you're interested, I'd love to provide more details. Looking forward to hearing from you.

D.3 Recruiters' Instructions

Please review the following information carefully, as it contains important details about the task at hand.

We are a specialized software company dedicated to crafting AI-driven solutions. As a venture-backed, cutting-edge firm, we are experiencing rapid growth and have numerous exciting openings across multiple locations. Currently, we're in a mass hiring period for computer software engineer roles in Chicago, Boston, Austin, and Seattle. Given the scale of our hiring needs, we're seeking your expertise to finalize compensation packages for several key positions within our tech firm.

Please review the attached job posting, which provides detailed information regarding the responsibilities and requirements for these positions. **Our allocated salary range for these roles is $115,000 to $155,000. However, we are open to considering offers outside this range for exceptional candidates. [V1 To adhere to local labor laws concerning pay transparency in certain offices, we have explicitly mentioned this range in our job advertisement.] [V2 We did not mention the target salary in the job advertisement.]**

We have extended informal offers to candidates, with an **initial offer of $130,000**. However, we have received indications from some of the candidates that they may be open to negotiation and have suggested counteroffers. With your insights, we aim to navigate these negotiations effectively, ensuring competitive compensation packages tailored to each candidate's unique

profile. All candidates have successfully cleared our rigorous technical assessments.

Given the different responses we received from the candidates, we are now seeking input from a compensation expert regarding suitable final compensation packages. We will share resumes of the 16 candidates who suggested counteroffers along with their respective counteroffer amounts. To streamline the process and gather your feedback efficiently, we have prepared a survey form, which we kindly request you to complete.

We estimate that your involvement in this task will require approximately 50 minutes of your time, followed by a brief 10-minute debriefing session. Your prompt response would be greatly appreciated, as we are eager to finalize our hiring decisions. Thank you for your time and contribution to our hiring efforts.

Please refrain from reaching out to any of these candidates. Your task solely involves evaluating them and providing us with your private assessments. Our team will handle further communication with the candidates. To avoid unintentional contact, we have redacted their contact details in the application forms provided. Our hiring strategy emphasizes making decisions exclusively on the submitted materials and interviews. Therefore, we advise against seeking additional information about individual candidates beyond the materials we provide, including online platforms.

Attached please find the following materials:

1. Information on the advertised openings
2. Resumes of the 16 candidates who suggested counteroffers
3. An Excel sheet detailing counteroffers proposed by the candidates
4. A link to a questionnaire to organize your suggestions: [LINK]

Your Contract ID is XX. For payment purposes, please indicate that when asked in the questionnaire.

D.4 Email – Task Description

Hello,

Thank you for your interest in the job post. I wanted to share a bit more information about the project with you.

As outlined, our company is currently experiencing rapid growth and has numerous exciting openings across multiple locations. Given the scale of our hiring efforts, we're seeking external assistance in candidate evaluation and finalizing compensation packages.

We will provide resumes of approximately 16 candidates in the late stages of recruitment. Your task will involve evaluating these candidate profiles and advising us on their compensation. To streamline this process and gather your feedback efficiently, we have prepared instructions and a survey form consisting of around 10 questions for each candidate. At this point, we do not need you to contact or seek additional information about the candidates themselves.

We estimate that your involvement in this task will require approximately 50 minutes of your time, followed by a brief 10-minute debriefing session. If you are interested, please let me know, and I will be in contact shortly. Please note that this is a one-time opportunity requiring one hour of work. If you believe you will require additional time to complete this task, please reach out in advance to request additional hours of work. Thank you for considering this opportunity.

D.5 Job Posting

Salary Range: $115,000 - $155,000 per year About Us: We're on the lookout for a Senior Software Engineer who's ready to dive into the dynamic world of technology solutions. Join our team and play a pivotal role in crafting and refining software applications that power our business operations. From collaborating with stakeholders to implementing innovative solutions, you'll be at the forefront of driving efficiency and excellence. Job Details: As a Senior Software Engineer, you'll be an integral part of our tight-knit team, dedicated to pushing the boundaries of innovation and service excellence. With a focus on cutting-edge technologies, you'll make meaningful contributions to our business and industry. Responsibilities:

- Design, develop, and implement top-tier software solutions aligned with user needs and business objectives.
- Spearhead the development of core platform capabilities for our innovative product.
- Foster collaboration among team members and stakeholders to cultivate a culture of teamwork and innovation.

- Lead all stages of the development lifecycle, from requirements analysis to production deployment.
- Drive continuous improvement efforts, enhancing software architecture and development processes.
- Ensure the performance, quality, and responsiveness of our applications.
- Develop APIs for seamless integration with other software components.
- Craft user-friendly UIs using modern technologies like React.
- Provide technical leadership and mentorship to junior team members, contributing to team growth.

Qualifications:

- 3-10 years of experience in software development.
- Proficiency in C#, .NET, React, and Django.
- Strong background in API and UI development.
- Deep understanding of the financial services industry and consumer finance.
- Proven track record of delivering high-quality software in fast-paced environments.
- Exceptional problem-solving skills with a creative and analytical mindset.
- Excellent communication and collaboration abilities, both independently and within a team.
- Bachelor's degree in Computer Science, Information Technology, or related field, or equivalent experience.
- Self-motivated with a drive to contribute to company success.
- Passion for technology and a commitment to staying abreast of industry trends.

If you're ready to make your mark and drive innovation alongside a talented team, we want to hear from you! Apply now and embark on a rewarding journey with us.

FIGURE C4. Example of Resumes

CBEB_5012

Benjamin [redacted]

Candidate Information

Candidate ID: CBEB_5012 **Mailing Address:** [redacted]
ZIP: [redacted] Phone : [redacted] **City/State:** [redacted]
Email: [redacted] **URL:** [redacted]
Are you legally authorized to work in the US? Y **Are you over the age of 18?** Y
Are you willing to relocate? Y **Will you require visa sponsorship?** N

Employment History (Last three jobs)

Title: Software Engineer
Company name: In-Depth Engineering
Location: Columbia, MD **Dates:** 2022-05 - Present

- Identify and patch software bugs within the AEGIS Combat System, utilizing C/C++, Java and Git for issue tracking and version control
- Develop full stack web applications that track shared engineering design documents
- Create and deploy REST endpoints that accept and serve engineering design drawings
- Migrate Access databases to MySQL, improving scalability and performance

Title: Software Engineer
Company name: Procraft, LLC
Location: Remote **Dates:** 2021-02 - 2022-05

- Managed and maintained the entire codebase for Soulveyor.
- Merged the Steamworks API throughout the code to allow the user's Steam client to interact with thegame.
- Designed a fully-fledged client prediction and server reconciliation framework entirely within C# toreliably ensure state across all clients, and to reduce RTT network latency by an average of 22ms.
- Created custom editor tools to reduce the asset/model/texture iteration, importation, and integrationworkload from over10 minutes of repetitive tasks, down to the simple click of a button.
- Orchestrated development for mobile devices, focusing on gameplay mechanics, performanceoptimizations, and ergonomicUI which resulted in a runtime performance improvement by an average of 54 FPS.

Title: Software Developer
Company name: Telesis
Location: Baltimore, MD **Dates:** 2018-09 - 2021-02

E Resume Construction

We developed a structured approach to resume construction. All resumes follow a standardized format, giving the impression that they are internally generated. Each resume is assigned a unique identification number within the material packet. The resumes were sourced from 208 actual resumes found on popular job board sites, focusing on candidates with software engineering experience from major US cities. Each resume includes three primary randomizing modules: work history, education, and first name.

The work history section features up to the last three jobs, with each job tenure ranging from 3 to 36 months. For the most recent position, there is a 50 percent chance of indicating current employment, and unemployment spells between roles range from 0 to 3 months. Each job entry combines the company name, location, and job description, which are drawn from the job experience section in the pool of sourced resumes. This means that our tailored resumes may combine elements, such as the first job from one version and the second job from another, to best match the job description. The education section includes institution names, attendance dates, and degrees, with each element independently selected from the resume pool to enhance variety. For names, first names are drawn from a list of popular baby names from the last 10 years, based on official government data (Social Security Administration), to ensure realism and cultural relevance.

Using this approach, we created 32 resume templates, from which we randomly selected 16 resumes to form a packet. To ensure realism and consistency, we take several additional steps. First, we use GEN-AI to paraphrase job descriptions within the job experience section to ensure that all three jobs maintain consistent writing styles. Second, we manually review all 32 resumes for consistency in terms of career progression and overall coherence within each template and across the entire packet. Third, we consulted professional recruiters to ensure that all materials align with industry standards.

F Questionnaire to Recruiters

F.1 First Questionnaire to Recruiters

We employ a systematic approach to recruitment, utilizing comprehensive questionnaires as a pivotal part of our screening process. This method ensures that we thoroughly assess candidates based on both their qualifications and their potential fit within our organization. Please choose one answer for each of the question below.

1. How would you evaluate each candidate in the following dimensions?

	Exceptional	Above Average	Below Average	Inadequate
Technical Skills				
Team skills or ability to work well with others				
Initiative				

2. What is your best estimate of the candidate's **current annual base salary** (in their most recent job)? Please provide a numerical number in dollars.

3. Based on your understanding of the candidate's experience, industry standards, and any information available about potential outside offers from competitors, what is your estimate for the candidate's **best annual outside offer**? Please provide a numerical number in dollars.

4. Given the candidate's skills, experience, and market standards, what **annual salary offer** would you recommend to remain competitive and attract the candidate? Please provide a specific numerical recommendation in dollars.

5. Our company offers varying levels of work-from-home days based on experience and qualifications. Given the candidate's preferences and the competitiveness of the offer, do you recommend including **work-from-home days** in the compensation package to enhance its appeal?"

   (a) Strongly recommend.

(b) Moderately recommend.

(c) Moderately advise against.

(d) Strongly advise against.

6. Given the recent developments at our company, a new, higher-level position may become available. Reflecting on your assessment of the candidate's qualifications and their negotiation behavior, would you recommend considering this individual for the **newly opened, more senior role**?

   (a) Strongly recommend.

   (b) Moderately recommend.

   (c) Moderately advise against.

   (d) Strongly advise against.

7. Do you suggest accepting the candidate's counteroffer?

   (a) Yes

   (b) No

8. If you answered 'No' in question 8, how would you recommend getting back to the candidate?

   (a) Attempt to negotiate with this candidate to reach closer to the initially offered salary.

   (b) Explore negotiating with this candidate by offering non-salary benefits but maintaining our initial offer.

   (c) Reject this application.

   (d) Other: please specify

9. Based on your assessment of the candidate's profile and our job opening, what is the likelihood that the candidate would accept the job offer if presented with the compensation package you recommended?

   (a) Very likely

(b) Likely

(c) Unlikely

(d) Very unlikely

F.2 Post-Experiment Questionnaire to Recruiters

As part of our debriefing process, we would like to ask some follow-up questions. To gather your feedback more effectively, we have prepared this survey form.

1. From your experience as a recruiter, how would you rate the overall quality of these 16 of candidates who applied to our company relative to other tech companies.

    - Excellent
    - Above average
    - Average
    - Below average

2. How informative are the candidates' counteroffers regarding their current salaries?

    - Very informative
    - Somewhat informative
    - Slightly informative
    - Not informative

3. How informative are the candidates' counteroffers regarding other opportunities the candidates have apart from the positions at our firm?

    - Very informative
    - Somewhat informative
    - Slightly informative
    - Not informative

4. How informative are the candidates' counteroffers regarding their personalities?

- Very informative
- Somewhat informative
- Slightly informative
- Not informative

5. Of the 16 candidates you reviewed, based on quality alone, which three candidates stand out the most in terms of qualifications? Please provide their candidate numbers and names.

6. In your professional experience, do you consider any of these reasons to be legitimate concerns that other firms may have when accommodating higher salary requests from candidates above the target pay range for the position? Feel free to select more than one or choose none of the above.

    - Fear of wider re-negotiations among employees or the necessity of offering higher salaries to future candidates
    - Considerations of fairness to other candidates regarding exceptions to the salary range
    - Concerns about potential legal consequences for making exceptions.
    - None of the above

7. In your experience, for job postings like ours, how common is it for candidates to suggest counteroffers exceeding the target pay range?

    - Very common
    - Common
    - Rare
    - Very rare

8. In your experience, for job postings like ours, how common is it for female candidates to suggest counteroffers exceeding the target pay range?

    - Very common

- Common
- Rare
- Very rare