Choice as Signal: Designing AI Adoption in Labor Market Screening

Brian Jabarian Pëllumb Reshidi

[Preliminary version]

December 3, 2025

Abstract

We study the design of human–AI screening systems in a hiring environment where applicants choose between interviewing with a human recruiter or an AI voice agent, a choice that creates a new and informative signal. Once firms condition on screener choice, welfare reversals emerge: choice benefits firms and high-ability applicants, but leaves low-ability ones worse off than under a predetermined assignment of the screening technology. Using data from a large-scale field experiment—where 70,000 applicants were randomly assigned to a human interviewer, an AI agent, or allowed to choose between them, we develop a structural estimation to quantify how choice-as-signal shapes match quality and welfare. This framework also allows us to evaluate alternative screening systems, with preliminary results suggesting higher welfare under joint human–AI screening than under either technology alone. Overall, we show that AI adoption in screening is a design problem rather than a simple human-versus-AI substitution decision.

Contact: Brian Jabarian (Corresponding author): Booth School of Business, University of Chicago, CDR, CAAI, CESifo, J-PAL, IGL, JILAEE, brian.jabarian@chicagobooth.edu. Pëllumb Reshidi; Department of Economics, Florida State University, preshidi@fsu.edu.

Acknowledgements: We are particularly grateful to Alex Imas for guidance. We also thank Luca Henkel, Eugenio Piga, Mariya Pominova, and Gregory Sun for valuable feedback and suggestions. Bernard Shen and Andrew James provided outstanding research assistance.

Partnership & Support: At PSG Global Solutions, we are deeply grateful to Vivek Padmanabhan, David Koch, David Bethea, and Anand Krishnan for their innovative leadership and unwavering scientific partnership. We also thank Glaiza Acevedo, Pritam Gawai, Rob McIntosh, Anand Narole, Jon-Ray Riveras, Ram Sabaru, James Shen, and Shraddha Bhujbal for their indispensable logistical support. At Google Cloud, we are grateful for the infrastructure and engineering support provided by Karim Atek, Brad Burns, and John Cecala. At Chicago Booth, we thank John Buenger and Michael Wu, Rich Cortez and Yvonne McConico, and Cheryl Danton for their guidance on legal, compliance, ethical, and data governance matters. Finally, we are grateful to Emily Bembeneck and Amy Boonstra for their administrative support.

Funding: Brian Jabarian gratefully acknowledges funding support from the University of Chicago Booth Center for Applied Artificial Intelligence, the Becker-Friedman Institute Program in Behavioral Economics Research, Google Cloud Research Program, Effective Venture US and the Swiss National Science Foundation Grant #10001768.

Transparency: The data used in this paper comes from Jabarian and Henkel (2025), study approved by the ethical committee of the University of Chicago (#IRB24-1894 and #IRB25-1002), collected through a field experiment pre-registered in the AEA RCT Registry (#15385) and under the Data-Use Agreement #160600008471, established by Brian Jabarian through the Booth School of Business at the University of Chicago, with PSG Global Solutions, a subsidiary of Teleperformance.

Disclosure: The partner firm supplied the data but had no role in analyses, manuscript preparation, or the decision to publish. Both authors and research assistants declare no financial or personal conflicts of interest during the conduct of this study. The data collection was completed between March 7 and June 7, 2025. On August 1, 2025, Brian Jabarian accepted an unpaid role as Chief Economist at PSG Global Solutions (a Teleperformance subsidiary), renewing the existing research partnership over the next five years, providing extended access for data and field experiments within the same pre-agreed publication rights as in the original DUA #160600008471.

1 Introduction

Screening technologies shape who gets hired, where talent goes, and how gains are shared between workers and firms. Employers translate noisy signals—such as credentials, tests, and interviews—into beliefs about applicants' ability, motivation, and fit, which, in turn, determine hiring decisions. Recent evidence from Jabarian and Henkel (2025) shows that AI voice agents perform well in the complex task of conducting live job interviews, improving worker–firm match quality relative to human-led interviews. The fact that AI agents can match—and in some cases exceed—human performance in live screening raises an important design question: should firms rely fully on AI? Is there still a reason to assign applicants to a human screener? Or perhaps a more flexible approach is for firms to offer applicants the option to choose between a human and an AI screener. These design choices have implications for the information firms collect, applicants' sorting behavior, and overall welfare.

In this paper, we develop a theoretical framework and an empirical structural analysis to study screener assignment as a design problem. We first model screener choice as a potentially informative signal in a multidimensional-skill environment with two screening technologies—human and AI. Using this model, we then analyze three information structures that govern how firms and applicants interpret choice as a signal. Finally, to estimate the model's key primitives, we leverage data from a large-scale field experiment, conducted in partnership with a recruiting firm, in which applicants were interviewed by a human recruiter, an AI voice agent, or given a choice between the two. The resulting estimates allow us both to quantify how screener choice affects hiring and sorting outcomes and to assess the welfare effects of alternative screening systems, including hybrid approaches that combine human and AI recruiters.

Our key contribution is to show that, when human and AI interviewers differ in how precisely they measure the various dimensions of applicants' multidimensional skills, *screener choice itself becomes an informative signal*. An applicant's decision to be interviewed by a human or an AI agent reveals private information about their strengths, weaknesses, or expectations about relative performance. Once screener choice is informative, screener assignment is no longer a mechanical step in the hiring process but an endogenous design choice that determines inference, sorting, and welfare.

Two mechanisms drive the welfare implications of screener assignment. *Comparative-advantage sorting* emerges because human and AI interviewers differ in how precisely they measure components of a multidimensional skill vector. When applicants can choose their screener, privately informed workers sort into the technology they expect to generate the

most favorable signals; when firms ignore screener choice, this ability to choose benefits workers while reducing firm surplus. A second mechanism, *inference from choice*, reflects that screener choice is made by applicants who know their private information: selecting a human versus an AI interviewer is itself an additional signal of type. Once firms condition on this choice, the additional inference reallocates surplus— benefiting firms and high-ability applicants while lowering the surplus of low-ability applicants. These two mechanisms operate differently across the three information structures we study.

We formalize these mechanisms in a model in which screeners generate noisy signals of multiple skill dimensions and may differ in how precisely they measure each dimension, and then analyze three information structures. In the first, the firm ignores screener choice entirely. In the second, the firm internalizes the information revealed by choice, but applicants do not anticipate this. In the third, both sides recognize that screener choice is an equilibrium signal that will be used in the hiring decision. These three structures generate distinct welfare predictions. When firms ignore choice, offering choice harms the firm and benefits all workers. When firms internalize choice but applicants do not, firms always gain, while low-ability workers are harmed and high-ability workers benefit. In the full equilibrium environment, firms typically weakly gain, high-ability applicants benefit, and low-ability applicants are again harmed. Thus, the welfare consequences of offering choice depend on how the informational content of choice is interpreted.

We complement the theoretical analysis with an empirical component that draws on a large-scale natural field experiment from Jabarian and Henkel (2025), conducted in partnership with a recruiting company, in which 70,000 applicants were randomly assigned to be interviewed by a human recruiter, an AI voice agent, or given the option to choose between them. This randomization provides clean variation in both the screening technology used to collect signals and the conditions under which the applicants reveal their screening preferences. Using this variation, we set up an estimation procedure that recovers key primitives of our model, including the relative precision of human and AI screeners across different skill dimensions, the firm's selectivity, and how the firm values each dimension when making hiring decisions. The procedure also allows us to distinguish among the three information structures by identifying whether firms internalize screener choice as a signal and, if so, whether applicants anticipate this.

We carried out a preliminary estimation–calibration exercise using the data from the deterministic human and deterministic AI treatments. These results allow us to conduct two preliminary counterfactual exercises. First, we assess the gains from giving the firm access to both human and AI interviews, rather than relying on a single source, and our preliminary findings point to substantial potential benefits from combining the two

information-collection technologies. Second, we evaluate the gains from assigning screeners to the skill dimensions they measure most accurately, illustrating the value of a hybrid screening-allocation strategy. We are currently completing the full estimation procedure, which will allow us to compare the choice regime with human-only and AI-only regimes to determine whether the information contained in applicants' selections can outperform deterministic screening assignments, and to study how alternative information structures shape equilibrium behavior and welfare.

Our analysis delivers two possibly counterintuitive policy implications. The first concerns choice. Although the choice branch is intended to grant applicants autonomy over screening technology, once firms can condition on screener choice, the act of selecting a human or an AI interviewer becomes an informative signal of underlying ability and, in equilibrium, tends to leave low-ability applicants worse off than under exogenous assignment. Allowing applicants to select their screener generally benefits firms and high-ability applicants, while low-ability applicants are typically worse off except in rare cases. If a policymaker's goal is to protect weaker candidates, exogenously assigning them to AI screening or limiting the extent to which firms may condition on screener choice can therefore dominate simple autonomy-based regimes. Substantively, this points to a "right not to be screened on one's choice" as a natural complement to a "right to choose," though enforcing such a right is challenging whenever screener choice is observable.

The second implication concerns automation. Even when AI interviewers match or exceed human performance, our estimates show that it can be optimal to keep humans and AI performing ostensibly the same task in systematically different ways. Welfare depends less on the absolute performance of each screener and more on how the screening system is designed—whether screeners are assigned or chosen, whether firms rely on a single or multiple interviewers, and how interview content and screener-choice signals are incorporated into decisions. In the hybrid screening systems we simulate, we employ both humans and AI, or have humans and AI specialize in different parts of the process, and show that these designs deliver higher firm surplus and better outcomes for selected groups of applicants than relying on either technology in isolation.

Taken together, our results suggest that AI automation in labor markets is less about one-to-one substitution of human labor and more about how human and AI work are jointly designed.

Related literature. Our paper contributes to three strands of literature: AI adoption in firms, human-AI collaboration in decision-making, and labor market screening.

AI adoption in firms. First, we relate to work on AI adoption in firms, which studies how AI reshapes production, labor demand, and organizational design (e.g., Brynjolfsson and McAfee 2014; Acemoglu and Restrepo 2018, 2019; Iansiti and Lakhani 2020; Czarnitzki et al. 2023; Dell'Acqua et al. 2023; Peng et al. 2023; Bonney et al. 2024; McElheran et al. 2024; Otis et al. 2024; Brynjolfsson et al. 2025); our analysis brings a comparable welfare and heterogeneity lens to hiring and screening. Our model of AI doption departs from principal–agent models in which a firm unilaterally decides whether to delegate a task to AI and then constrains human discretion (e.g., Athey et al. 2020), and instead treat adoption regimes as choice architectures that embed both paternalism and autonomy: purely human or purely AI screening corresponds to "technology paternalism" about how screening should be done, while a choice regime grants applicants autonomy but also generates an additional signal that firms may exploit. Our approach connects to the literature on paternalism, autonomy, and welfare in policy design (e.g., Bernheim and Rangel 2007, 2009; Bernheim 2016; Ambuehl et al. 2021) and extends it by considering "AI choice" as an adoption regime. Building on this regime, we then assess different human-AI screening systems—human-only, AI-only, or candidate choice— inducing a paternalism trade-off in which regulators and firms must balance the efficiency losses from adverse selection with the intrinsic value of respecting autonomy, consistent with evidence that preserving choice in technology adoption can increase long-term adherence to expert recommendations (Bushong et al. 2025).

Task-based and macro perspectives emphasize how automation reallocates tasks between humans and machines, and affects inequality within humans (e.g., Benzell et al. 2015; Brynjolfsson et al. 2018; Acemoglu and Restrepo 2022) By adopting a design-based approach to task modeling, our results show that automation need not take the form of one-for-one substitution of human labor by AI agents (a zero-sum game between humans and AI). Instead, it allows treating automation as an opportunity to redefine the purpose of tasks and expertise and to maximize gains by having both humans and AI agents perform the *same nominal task*—conducting job interviews—but with specialized roles, so that the total surplus created by screening expands rather than being merely redistributed between humans and AI (Autor and Thompson 2025).

Human–AI collaboration. Second, we contribute to the literature on human–AI collaboration, which studies how to combine algorithmic tools and human judgment in high-stakes decisions (e.g., Kleinberg et al. 2018). Prior work shows that AI assistance can be underweighted or misused, but that calibrated or coarsened risk scores can improve decisions relative to both raw algorithmic advice and human-only benchmarks (Agarwal

et al. 2023, 2025; Hoong and Dreyfuss 2025); human overrides of algorithmic recommendations often degrade performance (Stevenson and Doleac 2024), although a skilled minority can add value using private information (Hoffman et al. 2018a; Angelova et al. 2023), and theoretical work on delegation highlights the role of combining soft information with hard constraints (Frankel 2021). In hiring, hybrid human–AI models in which AI analyzes conversational interviews and humans evaluate them can improve workforce quality (Chakraborty et al. 2024), while work on selection pipelines cautions that intermediate screeners can destroy value if they are too correlated with the final assessor (Parasurama and Ipeirotis 2025); LLM-based interviews can capture richer textual information than simple tests and have matched or exceeded human-led interviews on technical question quality and conversational flow (Aka et al. 2025). Our results complement evidence that generative-AI tools deployed in production environments can raise average performance with substantial heterogeneity across workers and tasks (Brynjolfsson et al. 2025), by showing how alternative ways of embedding AI into screening—as a substitute, complement, or choice option—affect welfare and the distribution of gains.

Screening and labor market. Third, we connect to the broader literature on screening and selection in labor and industrial organizations. Previous work studies how interviews, tests, and other mechanisms affect labor market outcomes and organizational performance (e.g., McDaniel et al. 1994; Estrada 2019; Hoffman and Stanton 2024; Li et al. 2025; Radbruch and Schiprowski 2025). We build on this literature by modeling applicants with multidimensional skills and screening technologies with heterogeneous capabilities. As AI increasingly commoditizes traditional application materials and compresses variation in written signals (Cowgill et al. 2024; Wiles and Horton 2024; Galdin and Silbert 2025), the comparative value of live screening increases, whether conducted by humans or AI. Ash et al. (2025) shows that candidates' self-reported perceptions of interview difficulty predict hiring outcomes. In our setting, interviewer choice can reveal such a perception—so that screener assignment and screener choice jointly shape the information produced in interviews and the resulting welfare for both sides of the market.

We conceptualize applicants as possessing a vector of latent attributes imperfectly observed by the firm, in line with multidimensional signaling models of hiring (e.g., Vravosinos 2025), without assuming that AI or human screening dominates in any dimension. Field experiments show that AI can outperform humans in assessing "soft skills" (Cowgill 2019), interaction and communication skills (Jabarian and Henkel 2025) and technical skills (Aka et al. 2025). Our framework is also related to work that evaluates AI adoption using welfare rather than accuracy alone (e.g., Yang et al. 2023): we similarly take wel-

fare as the objective but focus on how screener assignment and screener choice reallocate surplus between firms and heterogeneous applicants once choice is itself an informative signal.

Paper structure. The remainder of the paper is organized as follows. Section 2 provides a background on the recruitment context and the institutional setting. Section 3 develops a screening technology choice model and presents results under Gaussian types and signals, beginning with the non-selection benchmark, moving to the case where firms incorporate applicant selection while workers are unaware, and finally analyzing the full equilibrium where both firms and applicants internalize selection. Section 4 sets out our estimation and calibration procedure, reports preliminary parameter estimates, and uses them to conduct counterfactual analyses of alternative AI adoption regimes and specialized screening designs. Section 5 concludes.

2 Institutional Setting and Data

2.1 Economic Environment

Firm partnership. We use de-identified data from the natural field experiment by Jabarian and Henkel (2025)[1], conducted in partnership with PSG Global Solutions, a Recruitment Process Outsourcing (RPO) firm integrated into Teleperformance, a large global Business Process Outsourcing (BPO) company. PSG specializes in high-volume recruitment for large U.S. and European clients in sectors such as telecommunications, finance, retail, technology, healthcare, and insurance, with recruiting centers worldwide. In the setting we study, PSG recruits entry-level customer-service representatives for call-center roles based in the Philippines. These positions pay substantially above the local minimum wage and require English proficiency, strong communication skills, willingness to work in changing shifts, and basic analytical and problem-solving skills. Appendix Section G provides an example of a detailed job description.

Industry background. The global Business Process Outsourcing (BPO) sector is a major employer and is characterized by very high turnover, forcing firms into continuous

[1]Beyond the administrative and test data we use here, the authors collected detailed interview transcripts and derived text measures, as well as applicant and recruiter surveys. These additional data are central to their analysis of mechanisms (e.g., "controlled variance"). Still, they are not required for our preliminary structural analysis of interviewer choice and adoption regimes, so we do not exploit them further in this current version. See Appendix J for more details on the IRB and pre-registration process of the experiment.

high-volume recruitment and making early-stage screening decisions first-order for both recruitment costs and worker–firm match quality. In the Philippines, BPO is one of the main hubs in the world for outsourced call-center work, with a large workforce serving predominantly U.S. and European clients (see, for example, Hernandez (2024)). On the supply side, a relatively young and well-educated labor force with strong English skills makes call-center jobs attractive relative to local alternatives; on the demand side, wages remain low compared to high-income countries and Filipino accents are perceived as close to standard U.S. English. These features, combined with very high turnover rates – on the order of 50–60 percent annually in call-center jobs – lead firms to invest heavily in continuous recruitment and training, making the design of early-stage screening particularly consequential (see, e.g., Berg et al. (2018), Sallaz (2019), and Buesing et al. (2020)).

2.2 Recruitment Process

Recruitment follows a multi-stage funnel, summarized in Figure F.1 in Appendix F. Applicants enter the pipeline either remotely or as walk-ins. Remote applicants pass a light engagement-based pre-screening, while eligible walk-ins are scheduled directly for an interview. The eligible applicants are then randomized into one of three screening systems (Human, AI voice agent, or Choice) and complete a 10- to 15 minute interview. Successful candidates take a standardized language and analytical test, after which a human recruiter reviews all the information and decides whether to extend a job offer. Applicants who accept proceed to training and early employment, during which we observe retention and, for a subset, on-the-job performance. The remainder of this subsection describes each stage of the funnel in more detail.

Recruitment modes and algorithmic pre-screening. In this setting, applicants learn about vacancies through standard channels: online job boards, social-media campaigns, the firm's website and digital campaigns, referrals, and walk-ins at recruitment sites. These channels generate two recruitment modes. In the *remote* mode, candidates submit a short online expression of interest with contact details and minimal background information. The firm computes an *engagement score* that aggregates signals such as completeness of the profile and responsiveness to outreach; applicants below an account-specific threshold are screened out, and the rest are queued for interview scheduling. PSG aims to interview as many candidates as possible and therefore sets low thresholds in the engagement-based pre-screening. See Appendix Section E.1 for details. In the *walk-in* mode, candidates present themselves in person at a recruitment site and, conditional on basic eligibility (e.g., age and legal work status), are directly scheduled for an interview without further pre-

screening. In both modes, invitations are sent via text messages and email, and an AI scheduler assigns interview slots to recruiters. See Appendix Section E.2 for the verbatim of the invitations.

Experimental randomization. Applicants who pass the engagement-based pre-screening are randomly assigned to one of three screening systems: (i) a *Human* regime, in which the interview is conducted by a human recruiter; (ii) an *AI* regime, in which the interview is conducted by an AI voice agent; and (iii) a *Choice* regime, in which applicants decide whether to be interviewed by a human recruiter or by the AI agent. Randomization occurs at the applicant level within the firm's recruitment system. In the Choice regime, the data record the actual interviewer chosen, so screener choice is an endogenous decision. Across all conditions, the data record whether the interview is completed, its timing, and the identity of the interviewer who actually conducted the interview.

AI voice agent. The AI interviewer, "Anna AI", is a cloud-based voice agent that conducts live job interviews over the phone. It combines standard conversational-AI components—automatic speech recognition, a large language model that selects questions and follow-ups, and text-to-speech synthesis—to implement the firm's structured interview in real time. We treat this AI agent as a given screening technology and focus on how its use should be designed relative to human interviewers. Appendix Section D provides additional technical details on the system architecture and implementation.

Job interviews. Both human recruiters and the AI agent ("Anna AI") follow the same structured interview guideline, which specifies key topics (including, among others, location and commute feasibility, schedule flexibility and attrition risks, motivation and career goals, prior work experience, education, and salary expectations), a recommended order, and example questions. A typical interview lasts 10–15 minutes, and recruiters retain flexibility to adjust the language and follow-up questions to the applicant's background. When the interview is conducted by Anna AI, the same guideline is implemented via the conversational agent. See Appendix Section H for two complete examples of interviews. In all conditions, a human recruiter evaluates the interview and makes hiring decisions; information collection (by the human or AI interviewer) and evaluation (by a human reviewer) are separate stages. At the start of the interview, the AI agent discloses its artificial identity and that a human recruiter will review the conversation and make the hiring decision.

Standardized language and analytical tests. Applicants who successfully complete the interview are invited to a standardized 30-minute test. The assessment consists of a lan-

guage module and an analytical module. The language module assesses English reading and writing skills and maps performance to the six CEFR proficiency bands (A1–C2). The analytical module aggregates three timed components - attention to detail, verbal reasoning, and numerical ability - into a composite score from 0 to 100. The test is taken after the interview and is required to advance to the hiring decision stage, so test scores are post-treatment and are not used as pre-treatment covariates in our main identification strategy.

Review and hiring decisions. Once both the interview and test results are available, a recruiter reviews the interview transcript and audio together with the standardized test scores and internal notes. Interviews are scored on a three-point scale based on customer service skills, communication skills, attrition risk, and salary fit, accompanied by a short free-text justification. Recruiters then decide whether the applicant is suitable for hire. Upon a positive general assessment and a suitable opening, the recruiter matches the applicant to a vacancy and issues a job offer; otherwise, the application is archived.

Downstream employment outcomes. The de-identified dataset we use links these recruitment records to downstream employment outcomes in the client firms. For each application, we observe the completion of the interview, the pass/fail decision at the interview stage, whether an offer is made, and whether the offer is accepted. For applicants who start work, we observe completion of the one-week training period and retention at several horizons (e.g., one, two, three, and four months). Because match quality is not directly observable, we follow the labor and personnel literature in using retention as a proxy for worker–firm match quality in this high-turnover environment (see, e.g., (Jovanovic 1979; Autor and Scarborough 2008; Hoffman et al. 2018b)). For a subset of workers, we also observe standard on-the-job productivity metrics, including quality-assurance scores (QA), average handling time (AHT), and customer-satisfaction ratings (CSAT). Data are organized at the applicant–application level and include pre-treatment characteristics (recruitment mode, engagement measures, posting and account), the assigned regime and realized interviewer, interview and test outcomes, and the sequence of hiring and early-retention outcomes.

2.3 Experimental Sample

Scope and timing. The underlying pre-registered natural field experiment, described in detail by Jabarian and Henkel (2025), was implemented between March 7 and June 7, 2025. During this period, PSG received 70,884 applications for 48 distinct entry-level

customer-service postings across 43 client accounts, including 23 Fortune 500 and 20 European leading companies in technology, insurance, telecommunications, retail, finance, healthcare, and transportation. Appendix Table I.1 reports the distribution of applications by client industry. Applications were processed by 26 recruitment sites located in 19 cities across several regions of the Philippines, with the bulk of the volume coming from Metro Manila and the Central Visayas region. See Appendix Figure F.2 for the distribution of applications per region.

Randomized sample. Applicants enter the pipeline through the two recruitment modes described above. Of the 70,884 applications, 67,056 meet the firm's eligibility criteria and are randomized into one of the three screening systems (Human, AI, or Choice). This randomized sample serves as the basis for our analysis. In total, 17,621 applications arrive through the walk-in mode and 53,263 through the remote mode. Walk-in applicants who meet the basic eligibility criteria are scheduled for an interview. In contrast, remote applicants are invited, subject to clearing a low threshold for the hiring process engagement, and are screened algorithmically.

Randomization yields 40,103 applications in the AI condition (10,421 walk-in; 29,682 remote), 13,557 applications in the Human condition (3,478 walk-in; 10,079 remote), and 13,396 applications in the Choice condition (3,469 walk-in; 9,927 remote). Jabarian and Henkel (2025) show that pre-treatment characteristics are balanced across these conditions (see their balance Table, reproduced in Appendix Table I.2). In total, 6,319 applications receive a job offer, and 4,294 can be linked to downstream employment records; we use this matched subset when we analyze retention and employment outcomes. Most applications are submitted by unique individuals: 64,556 applications correspond to distinct persons, so roughly 9% of applicants submit more than one application during the experimental window. In our structural analysis, we treat applications as the unit of observation and show in robustness exercises that excluding repeat applicants does not materially affect the results.

The data structure allows us to compare the informativeness of human and AI interviews, analyze screener choice in the choice condition, and calibrate counterfactual screening systems. To interpret these screening systems and evaluate the welfare effects of alternative policies, we now describe a model of screening technology choice.

3 A Model of Screening Technology Choice

We study a hiring environment in which applicants differ along multiple unobserved dimensions of ability and face a choice between alternative *screening technologies*. These technologies—such as human or AI interviewer—differ in how precisely they measure each relevant dimension of applicant quality. Human screeners, for example, may better assess communication skills or interpersonal fit but be noisier in evaluating technical aptitude, whereas AI systems may excel at analyzing structured responses or problem-solving tasks but perform less reliably on subjective or contextual traits. The decision of which screener to select thus represents a form of *self-selection*: applicants strategically choose the evaluation method under which they expect to appear stronger.

The firm, in turn, aims to infer overall productivity from the information generated by these heterogeneous screening technologies.[2] Depending on institutional design, the firm may or may not observe the applicant's screener choice, and may or may not take it into account when making hiring decisions. Our analysis builds from a baseline in which the firm does not internalize the screener choice to a fully strategic setting in which both the firm and applicants internalize the informational and signaling implications of that choice.

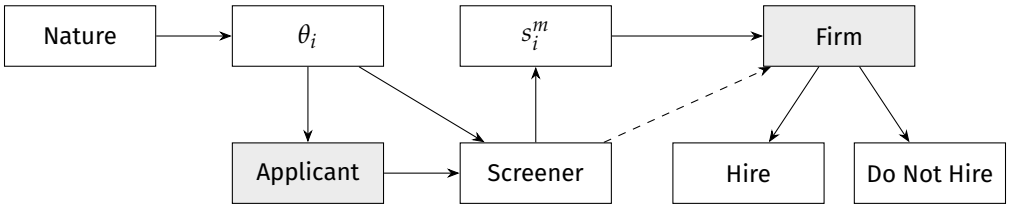
Figure 1 summarizes the sequence of events. Nature draws an applicant's underlying type, the applicant chooses a screening technology, the screener generates noisy signals about the applicant's attributes, and the firm observes the resulting evidence—possibly along with the information of whether the applicant chose the screener—before deciding whether to hire. This setup provides a general framework for studying how self-selection and inference may interact in modern hiring environments that combine human and AI-based assessment tools.

Below, we formalize the environment: while the type and signal structure can be defined in general terms, we focus on the case in which both follow Gaussian distributions, which provides analytical tractability.

In Section 3.1 we examine the benchmark case in which the candidate may choose between a human or AI screener, but the firm does not take this choice into account when making hiring decisions. Whether firms actually behave in this way is ultimately a mat-

[2] In line with standard terminology, we say that each technology $m \in \{Human, AI\}$ "generates" a signal s_j^m about applicant's j dimension. In our empirical setting, however, the AI interviewer does not produce a recommendation or hiring score. Instead, the AI interviewer generate a signal by conducting the job interview, i.e., *collecting* information from the applicant's answers (a transcript and derived features), which is then passed to the human recruiter. Throughout, references to "AI that generates a signal" should therefore be interpreted as the AI voice agent that generates and records informational content from the interview, not as AI that generates a recommendation or a final decision.

Figure 1: Flowchart of the Hiring Process



Notes: Nature draws the applicant's type θ_i. After observing their own type, the applicant selects a screener $m \in \{Human, AI\}$. The chosen screener then produces signals s_i^m, which the firm observes—potentially together with information on whether the applicant chose the screener or was preassigned to it—before making its hiring decision.

ter of institutional detail and may vary across settings.[3] Beginning with this environment is useful as it provides a baseline that reveals who stands to benefit and who is disadvantaged when such a choice is made available to candidates. This step also allows us to isolate the mechanical consequences of selection without yet introducing strategic responses on the part of either firms or workers.

We then turn to Section 3.2, where we consider the intermediate case in which firms do incorporate information about the candidate's choice of screener into their hiring decisions. In this setting, however, workers remain non-strategic: they do not anticipate that their choice of screener will itself be evaluated.[4] Comparing this case with the benchmark highlights how firm-side use of selection information alters outcomes, even before workers adjust their behavior in response.

Finally, in Section 3.3 we analyze the equilibrium outcomes in which both sides are aware that screener selection is part of the evaluation criterion. Here, driven by tractability and motivated by institutional details, we make a simplifying assumption that, after updating beliefs, the adjustments firms make to account for the chosen screener are independent of the realized signals. Firms explicitly condition their decisions on the candidate's choice, and applicants take this into account when deciding which screener to select. This final step closes the loop, capturing the strategic interaction that arises once both firms and applicants internalize the role of selection.

By structuring the analysis in this way, moving from the simplest benchmark to the fully strategic environment, we can disentangle the distinct forces at play. The comparison

[3]For example, if the hiring decisions are being made by workers who are not informed that applicants had a choice, then they would take the outcomes of the interviews at face value.

[4]This case is relevant if applicants believe their choice of screener will not enter the evaluation process. It may also arise if applicants lack the sophistication to recognize that their screener choice could itself be factored into the firm's hiring decision.

across these environments makes transparent which features of the choice benefit one side, which disadvantage the other, and under what conditions both parties may gain.

Types, Signals, and Selection Criteria Although the firm screens multiple candidates, the analysis can be carried out separately for each one. We therefore drop the candidate index i in this section and treat the problem as if there were a single applicant. We assume that both types and signals follow Gaussian distributions. For each dimension $j \in \mathcal{J}$, a candidate has an underlying type θ_j, while the chosen screener $m \in \mathcal{M}$ produces a noisy signal s_j^m about that type. Signals are drawn independently across j and across screeners m. The following assumption summarizes this structure formally.

Assumption 1 (Normality)**.** *Candidate types and signals are normally distributed. Specifically, for each dimension $j \in \{1, 2\}$,*

$$\theta_j \sim \mathcal{N}\left(\mu_j, \frac{1}{\rho_j}\right), \qquad s_j^m = \theta_j + \varepsilon_j^m, \quad \varepsilon_j^m \sim \mathcal{N}\left(0, \frac{1}{\rho_j^m}\right).$$

Here, ρ_j denotes the prior precision regarding θ_j, whereas ρ_j^m denotes the precision of the signal generated by screener m for dimension j.

For clarity, in the rest of our analysis we focus on the two-dimensional case $j \in 1, 2$. In practice, candidates may possess many characteristics, some of which are better assessed by human screeners and others by AI. Because the firm's evaluation criterion is linear, as will be described immediately below, it is without loss of generality to aggregate all characteristics evaluated with greater accuracy by humans into a single composite dimension, denoted θ_1, and those better evaluated by AI into another, denoted θ_2. Since the underlying characteristics are normally distributed, these aggregates will also be normally distributed. Hence, restricting attention to two dimensions suffices for our purposes while keeping the analysis tractable and clear.

After observing screener signals and updating beliefs, we assume the following:

Assumption 2 (Linear Aggregator and Hiring Criteria)**.** *The firm evaluates candidates using a linear index of posterior means:*

$$\hat{I}^m = \alpha \hat{\mu}_1^m + (1 - \alpha)\hat{\mu}_2^m, \qquad 0 \leq \alpha \leq 1,$$

and hires if and only if $\hat{I}^m \geq \bar{I}$ for some cutoff $\bar{I}$.

Here, $\hat{\mu}_1^m$ and $\hat{\mu}_2^m$ denote the posterior means. The parameter α reflects the relative importance the firm assigns to the first dimension, with $1 - \alpha$ the weight on the second,

so the index $\hat{I}^m$ provides a single summary measure of the candidate's expected value. This threshold rule reflects the idea that firms must ultimately make a binary decision, either to extend an offer or not, based on potentially complex information. Reducing the posterior estimates across dimensions to a weighted index provides a convenient way to place candidates on a common scale. Assuming that the decision is governed by a cutoff $\bar{I}$ is common in models of hiring and admissions, as it can capture capacity or quality constraints (e.g., only candidates above a certain expected productivity are worth hiring) and provide consistent treatment across applicants. In practice, firms often adopt similar procedures, for example, by assigning scores across multiple criteria and accepting those who exceed a benchmark. Although stylized, the rule offers analytical tractability and a plausible approximation of real-world decision processes.

In what follows, we also make two technical assumptions.

Assumption 3 (Technical Assumptions)**.**

Selective Threshold. *For each $j \in \{1, 2\}$, the prior mean lies below the hiring cutoff,*

$$\mu_j < \bar{I}.$$

Informative Screening. *For each $j \in \{1, 2\}$, the precision of the screener exceeds that of the prior,*

$$\rho_j^m > \rho_j.$$

The first technical assumption ensures that the unconditional distribution of applicant quality does not already exceed the hiring bar. Without such a restriction, the firm would find it profitable to hire the vast majority of the population without screening, which is inconsistent with selective industries in which only a fraction of candidates receive offers. The second captures the idea that the screening process is quite informative: interviews, documentation, and other background checks provide much richer information about an applicant's suitability than what can be inferred from prior population distribution alone. While these assumptions help keep the analysis concise and easy to present, they are not critical for our main results.

3.1 The No-Selection Benchmark

3.1.1 Firm's Inference and Hiring Decision

We now examine a benchmark environment in which candidates may opt for either a human or an AI screener, while the firm's hiring decision is independent of that choice.

Suppose an applicant chooses screener m, and the firm observes the two-dimensional signal vector

$$s^m = (s_1^m, s_2^m) \in \mathcal{S}_1 \times \mathcal{S}_2.$$

The firm then updates its belief about the candidate's type according to Bayes' rule:

$$\pi^m(\theta \mid s^m) = \frac{g^m(s^m \mid \theta)\, f(\theta)}{\int_\Theta g^m(s^m \mid \tilde{\theta})\, f(\tilde{\theta})\, d\tilde{\theta}},$$

where

$$f(\theta) = f_1(\theta_1) f_2(\theta_2)$$

denotes the prior density over types. Under the assumption of conditional independence across the two dimensions, the Gaussian case yields

$$\theta_j \mid s_j^m \sim \mathcal{N}\Bigg(\underbrace{\frac{\rho_j\, \mu_j + \rho_j^m\, s_j^m}{\rho_j + \rho_j^m}}_{\hat{\mu}_j^m(s_j^m)}, \underbrace{\frac{1}{\rho_j + \rho_j^m}}_{\hat{\rho}_j^m} \Bigg), \quad j \in \{1,2\}. \tag{1}$$

In each dimension, the posterior mean is a precision-weighted average of the prior mean and the observed signal, while the posterior variance decreases with signal precision. After observing the signal vector $s^m = (s_1^m, s_2^m)$, the firm forms posterior means for each dimension,

$$\hat{\mu}_j^m = E[\theta_j \mid s_j^m], \quad j \in \{1,2\}.$$

The firm then, as described earlier, constructs the linear index

$$\hat{I}^m = \alpha\, \hat{\mu}_1^m + (1-\alpha)\, \hat{\mu}_2^m, \quad \text{with } 0 \le \alpha \le 1, \tag{2}$$

and hires if and only if $\hat{I}^m \ge \bar{I}$ for some cutoff $\bar{I}$.

3.1.2 Applicant's Problem

Probability of Hire From the applicant's perspective, the key objects are the posterior means $\hat{\mu}_1^m$ and $\hat{\mu}_2^m$ that the firm forms about them after observing signals. Importantly, the applicant does not know in advance what signals screener m will generate, and therefore cannot know exactly what the firm's posterior will be. Instead, they must think in terms of the distribution of these posterior means.

Because signals are Gaussian with prior precision ρ_j and screener-m precision ρ_j^m, the

distribution of the posterior mean in each dimension is itself Gaussian:

$$\hat{\mu}_j^m \sim \mathcal{N}\Big(\tilde{\mu}_j^m(\theta_j),\ \tilde{\rho}_j^m\Big), \quad j \in \{1,2\},$$

where

$$\tilde{\mu}_j^m(\theta_j) = \frac{\rho_j\,\mu_j + \rho_j^m\,\theta_j}{\rho_j + \rho_j^m}, \qquad \tilde{\rho}_j^m = \frac{\rho_j^m}{(\rho_j + \rho_j^m)^2}.$$

Posterior means are independent across the two dimensions, so the applicant effectively faces two independent sources of noisy assessment. Since the firm combines the two dimensions into a single index, the applicant anticipates that the composite index under screener m will follow:

$$\hat{I}^m = \alpha\,\hat{\mu}_1^m + (1-\alpha)\,\hat{\mu}_2^m \sim \mathcal{N}\Big(\alpha\,\tilde{\mu}_1^m + (1-\alpha)\,\tilde{\mu}_2^m,\ \alpha^2\,\tilde{\rho}_1^m + (1-\alpha)^2\,\tilde{\rho}_2^m\Big).$$

In words, the applicant expects the firm's evaluation to be a weighted average of the two posterior means. The mean of this distribution reflects the expected score the applicant believes he will receive, while the variance captures the residual uncertainty due to signal randomness. Knowing their own type $\theta = (\theta_1, \theta_2)$, the applicant can then compute the probability of being hired under screener m:

$$p_m(\theta) = \Pr\big(\hat{I}^m \geq \bar{I} \mid \theta\big) = 1 - \Phi\left(\frac{\bar{I} - \big(\alpha\,\tilde{\mu}_1^m + (1-\alpha)\,\tilde{\mu}_2^m\big)}{\sqrt{\alpha^2\,\tilde{\rho}_1^m + (1-\alpha)^2\,\tilde{\rho}_2^m}}\right).$$

This expression makes clear how both the expected index and its uncertainty affect hiring chances. The numerator, $\bar{I} - (\alpha\,\tilde{\mu}_1^m + (1-\alpha)\,\tilde{\mu}_2^m)$, measures how far the cutoff $\bar{I}$ lies above or below what the applicant expects to score. The denominator, $\sqrt{\alpha^2\,\tilde{\rho}_1^m + (1-\alpha)^2\,\tilde{\rho}_2^m}$, rescales this difference by the variability of the index.

Naturally, a higher expected score always raises the applicant's probability of being hired: the larger their expected index, the more frequently they clear the hiring threshold. The role of index variability, however, is more subtle. If the applicant's expected score is above (below) the cutoff $\bar{I}$, the hiring probability is greater (less) than one-half. As the variability of the index grows, $p_m(\theta)$ converges to $\frac{1}{2}$ regardless of how high or low the expected score is. Thus, greater variability benefits applicants who expect to fall short of the cutoff, since it gives them a chance to be hired despite a low mean, but it harms those who expect to exceed the cutoff, since it introduces the risk of falling below it. Put differently, weaker applicants prefer noisier signals, while stronger applicants favor more precise signals that allow the firm to recognize their quality.

Applicants' Choice Each applicant chooses the screener that maximizes their probability of being hired:

$$m^*(\theta) = \arg\max_{m \in \mathcal{M}} \; p_m(\theta),$$

where $p_m(\theta)$ is the hiring probability under screener m given type $\theta = (\theta_1, \theta_2)$. Recall

$$p_m(\theta) = \Pr\big(\hat{I}^m \geq \bar{I} \mid \theta\big) = 1 - \Phi\left(\frac{\bar{I} - \left(\alpha \, \frac{\rho_1 \, \mu_1 + \rho_1^m \, \theta_1}{\rho_1 + \rho_1^m} + (1-\alpha) \, \frac{\rho_2 \, \mu_2 + \rho_2^m \, \theta_2}{\rho_2 + \rho_2^m} \right)}{\sqrt{\alpha^2 \, \frac{\rho_1^m}{(\rho_1 + \rho_1^m)^2} + (1-\alpha)^2 \, \frac{\rho_2^m}{(\rho_2 + \rho_2^m)^2}}} \right).$$

Since $\Phi(\cdot)$ is strictly increasing, the applicant compares the arguments of the normal CDF across screeners and chooses the one with the lower value. To characterize the indifference condition, for $m \in \{AI, H\}$ define

$$R_m \;=\; \sqrt{\left(\frac{\alpha^2 \, \rho_1^m}{(\rho_1 + \rho_1^m)^2} + \frac{(1-\alpha)^2 \, \rho_2^m}{(\rho_2 + \rho_2^m)^2} \right)}, \;\; A_m = \frac{\alpha \, \rho_1^m}{(\rho_1 + \rho_1^m) \, R_m}, \;\; B_m = \frac{(1-\alpha) \, \rho_2^m}{(\rho_2 + \rho_2^m) \, R_m}. \quad (3)$$

The applicant is indifferent between AI and H along a linear cutoff in (θ_1, θ_2)-space:

$$\theta_1 = \beta_0 + \beta_1 \theta_2,$$

with coefficients

$$\beta_0 \;=\; \frac{\bar{I} \, (R_{AI}^{-1} - R_H^{-1}) - \rho_1 \mu_1 \left(\frac{A_{AI}}{\rho_1^{AI}} - \frac{A_H}{\rho_1^H} \right) + \rho_2 \mu_2 \left(\frac{B_H}{\rho_2^H} - \frac{B_{AI}}{\rho_2^{AI}} \right)}{A_{AI} - A_H}, \quad \beta_1 \;=\; \frac{B_H - B_{AI}}{A_{AI} - A_H}. \quad (4)$$

Thus, the optimal screener choice is determined by which side of this cutoff line the applicant's type lies:

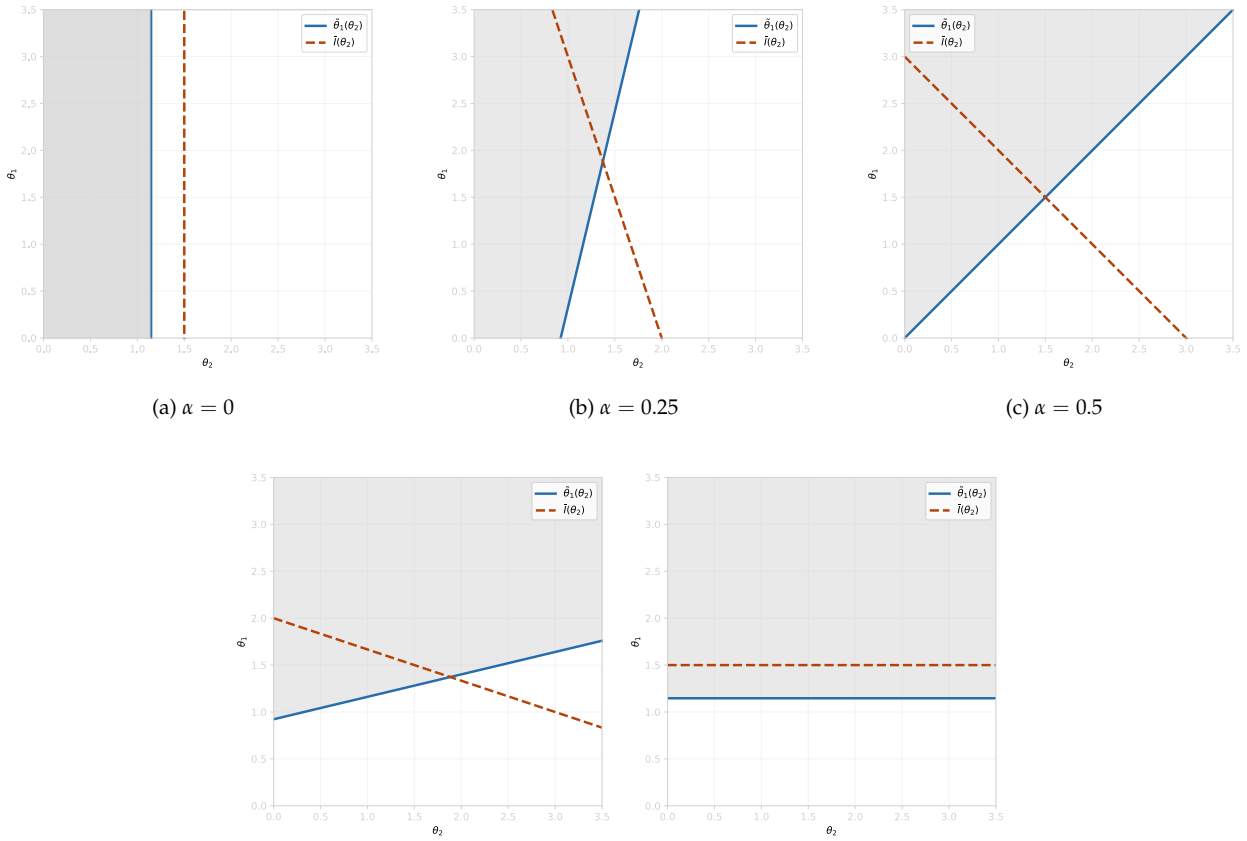
$$m^*(\theta) = \begin{cases} \text{AI}, & \text{if } \big(A_{AI} \geq A_H \text{ and } \theta_1 \geq \beta_0 + \beta_1 \theta_2\big) \;\text{ or }\; \big(A_{AI} < A_H \text{ and } \theta_1 \leq \beta_0 + \beta_1 \theta_2\big), \\ H, & \text{otherwise.} \end{cases}$$

In words, the applicant chooses AI whenever her type (θ_1, θ_2) lies on the AI side of the cutoff line, and H otherwise. The sign of $A_{AI} - A_H$ determines whether being "above" or "below" the line corresponds to choosing AI.

The intuition is as follows: although several parameters influence the cutoff, when one screener is substantially more precise in the first dimension, and the other is substantially

more precise in the second, the coefficient β_1 is typically positive. In such cases, the applicant's choice is determined by which attribute θ_1 or θ_2 is stronger, and naturally selects the screener that best reveals it. Figure 2 depicts such a case for different values of α. The figure shows the regions of the (θ_1, θ_2)-space where the applicants choose H (gray region) or AI (white region), together with the firm's recruitment cutoff $\bar{I}$ (dashed line). In all all examples shown, the human screener has a higher precision in θ_1 (the vertical dimension). Thus, when θ_1 is sufficiently large relative to θ_2, the applicant selects the human screener, whereas when θ_2 is relatively large compared to θ_1, the applicant selects the AI screener.

Figure 2: Applicants Optimal Choice: Undominated Screeners



(a) $\alpha = 0$ (b) $\alpha = 0.25$ (c) $\alpha = 0.5$

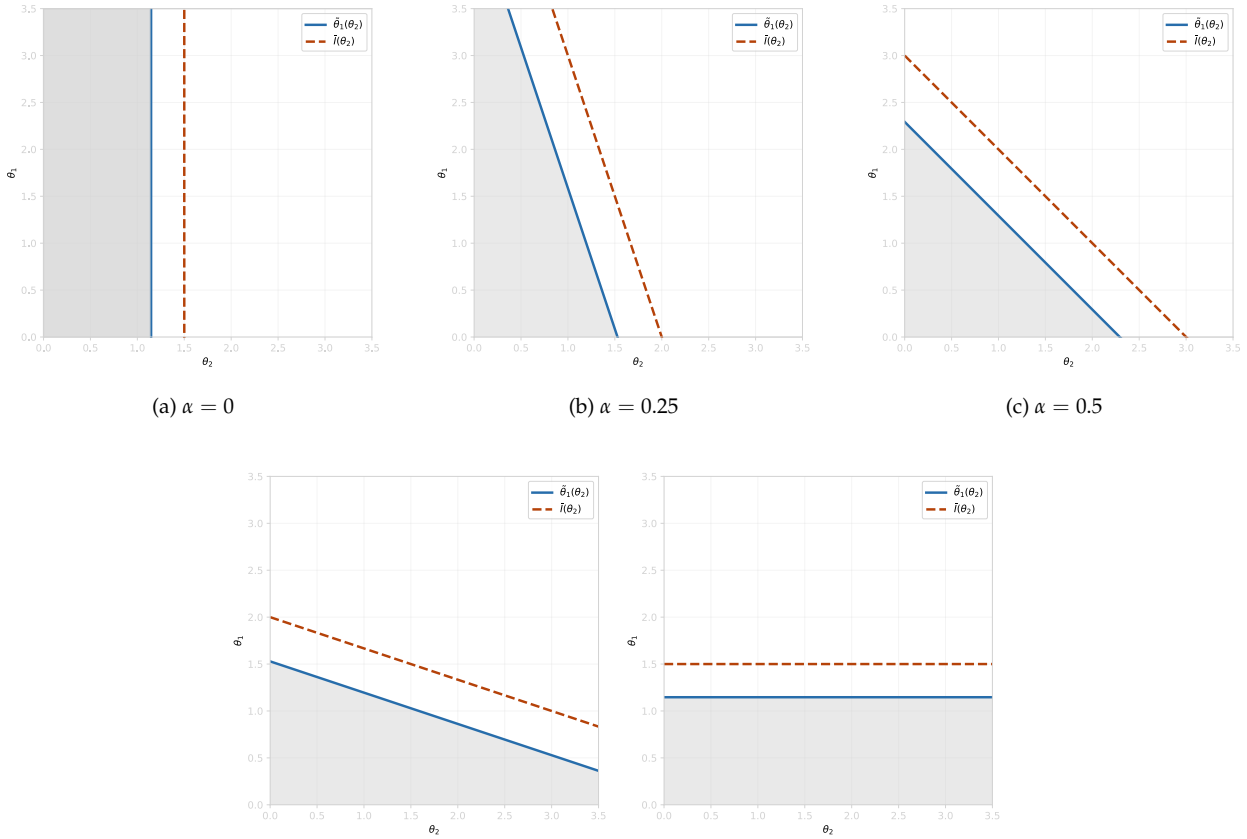
(d) $\alpha = 0.75$ (e) $\alpha = 1$

Notes: The blue cutoff represents the set of (θ_1, θ_2) combinations at which the candidate is indifferent between selecting the AI screener and the human (H) screener. We assume that $\rho_1^H > \rho_1^{AI}$ and $\rho_2^H < \rho_2^{AI}$—that is, the human is more precise in evaluating the first dimension, while the AI is more precise on the second. Under this assumption, the candidate prefers the human screener in the region above and to the left of the blue cutoff (the gray area), and the AI screener in the region below and to the right. The dashed orange cutoff indicates the hiring threshold—the level that must be met for the candidate to be hired. The parameter values used above are $\rho_1^H = 2$, $\rho_1^{AI} = 1$, $\rho_2^H = 1$, $\rho_2^{AI} = 2$, $\rho_1 = 1$, $\rho_2 = 1$, $\mu_1 = 1$,$\mu_2 = 1$, $\bar{I} = 1.5$, and $\alpha = \{0, 0.25, 0.50, 0.75, 1\}$.

When one screener dominates, meaning it is more precise on both dimensions, the

applicant's decision depends on her position relative to the hiring cutoff rather than on relative strengths in θ_1 or θ_2. Applicants who are just below or comfortably above $\bar{I}$, prefer the more precise screener, since it accurately reveals their high quality. However, those who are far below the cutoff prefer the noisier screener, hoping that favorable noise might push them above the threshold. Figure 3 illustrates this pattern: the dominant screener (AI, white region) is chosen in regions near and above the cutoff point (dashed line), while the dominated screener (noisier H, gray region) is chosen only by applicants sufficiently below it.

Figure 3: Applicants Optimal Choice: Dominated Screeners



(a) $\alpha = 0$ (b) $\alpha = 0.25$ (c) $\alpha = 0.5$

(d) $\alpha = 0.75$ (e) $\alpha = 1$

Notes: The blue cutoff represents the set of (θ_1, θ_2) combinations at which the candidate is indifferent between selecting the AI screener and the human (H) screener. We assume that $\rho_1^H < \rho_1^{AI}$ and $\rho_2^H < \rho_2^{AI}$—that is, the AI is more precise in evaluating both dimensions. Under this assumption, the candidate prefers the human screener in the region to the left and below the blue cutoff (the gray area), and the AI screener in the region to the right and below it. The dashed orange cutoff indicates the hiring threshold—the level that must be met for the candidate to be hired. The parameter values used above are $\rho_1^H = 1$, $\rho_1^{AI} = 2$, $\rho_2^H = 1$, $\rho_2^{AI} = 2$, $\rho_1 = 1$, $\rho_2 = 1$, $\mu_1 = 1$, $\mu_2 = 1$, $\bar{I} = 1.5$, and $\alpha = \{0, 0.25, 0.50, 0.75, 1\}$.

Observation 1 (Sign of the cutoff slope). *Fix two screeners $m, m' \in \{AI, H\}$.*

(i) Dominated screener. *For any $\alpha \in (0, 1)$, if $\rho_1^m < \rho_1^{m'}$ and $\rho_2^m < \rho_2^{m'}$ (i.e., m' is strictly more*

precise on both dimensions), then $\beta_1 < 0$.

(ii) Cross precision. *For any $\alpha \in (0,1)$, if at least one of the two inequalities in (i) fails – e.g., $\rho_1^{m'} > \rho_1^m$ but $\rho_2^{m'} < \rho_2^m$ (or the reverse) – then β_1 may be positive or negative.*

Part (i) covers the dominated case: when one screener is strictly more precise on both dimensions, the applicant knows that this screener always reveals her type more accurately. In this case, the indifference line must slope downward ($\beta_1 < 0$), meaning that applicants above the hiring bar prefer the more precise screener to signal their strength. In contrast, those far below prefer the noisier screener in the hope that randomness may push them above the threshold. Part (ii) concerns the cross-precision case, where one screener is more precise on θ_1 (say, H) while the other is more precise on θ_2 (say, AI). In this case, the sign of the cutoff slope β is parameter dependent (including α and the precision levels). Importantly, the fact that neither screener dominates the other does not imply that β must be positively sloped. To see this, consider cases in which the firm places an extremely high weight on the first dimension (i.e., α is very large) and is therefore only weakly concerned with the second. From the firm's perspective, it is as if the human screener dominates the AI screener, even though the AI screener produces more precise signals for the second dimension. In turn, the applicant continues to adopt a downward-sloping cutoff.[5]

3.1.3 The Welfare Consequences of Screener Choice

We now turn to the *welfare consequences* of allowing applicants to choose their interviewer, rather than being assigned to either a human or an AI screener. Throughout the paper, we evaluate welfare using simulations. Although closed-form expressions can be derived for the present section, they quickly become intractable in later analyses; for consistency and clarity, we therefore rely on simulation-based methods throughout. Our simulation design varies the precision of one screener over a wide range while holding the other screener fixed, allowing us to examine settings in which one screener dominates, in which

[5] Note that in both cases, as $\alpha \to 1$ the slope converges to 0, meaning the firm screens only on θ_1, while as $\alpha \to 0$ the slope converges to 1, so the firm screens only on θ_2. The way the slope evolves between these extremes depends on the screeners' relative precision and on whether we are in case (i) or (ii). In the first case, the slope is always negative between the points, whereas in the second case, the behavior depends on the parameter values. Further note that

$$\lim_{\alpha \to 1} \beta_0 + \beta_1 \theta_2 = \bar{I} - \frac{(\bar{I} - \mu_1)\rho_1}{\sqrt{\rho_1^{AI}\rho_2^H}}$$

That is, if the firm only cares about one component, the cutoff reduces to that in the one-dimensional case, see Appendix Section A.2.1.

both are similarly informative, and in which each is effectively specialized in a different dimension. With this framework in place, our central question is whether introducing choice improves outcomes for both sides of the market: for applicants, by increasing their chance of employment, and for firms, by improving match quality.

Firm Welfare To make this comparison precise, we begin by formalizing firm welfare. We define firm welfare as

$$W_F = \mathbb{1}\{\hat{I}^m > \bar{I}\}\left(\nu\,\mathbb{1}\{I \geq \bar{I}\} - \eta\,\mathbb{1}\{I < \bar{I}\}\right), \quad \nu, \eta > 0.$$

This expression can be unpacked into several intuitive components. The *outer indicator* term, $\mathbb{1}\{\hat{I}^m > \bar{I}\}$, governs the hiring decision itself. An applicant is only invited if their *expected* composite index $\hat{I}^m$ clears the firm's employment cutoff $\bar{I}$. This expectation reflects how the firm interprets the applicant's signals:

$$\hat{I}^m = \mathbb{E}[\alpha\theta_1 + (1-\alpha)\theta_2 \mid s_1^m, s_2^m] = \alpha\,\hat{\mu}_1^m + (1-\alpha)\,\hat{\mu}_2^m,$$

where $\hat{\mu}_1^m$ and $\hat{\mu}_2^m$ denote the posterior means of the applicant's type along each relevant dimension. Conditional on this employment, firm welfare then depends on the applicant's *true composite index*,

$$I = \alpha\theta_1 + (1-\alpha)\theta_2.$$

If the true index exceeds the cutoff $\bar{I}$, the hire is productive and generates a benefit of size ν. In contrast, if the true index falls below the cutoff, the hire is a mistake from the firm's perspective. It generates a cost of size η.[6] Taken together, this formulation captures two layers of the problem: the *ex ante* decision based on noisy signals and the *ex post* realization of the applicant's underlying productivity. It allows us to weigh the benefits of flexibility and choice in the interview stage against the risks of erroneous hiring, and ultimately to compare welfare across different institutional arrangements—choice, Human-only assignment, or AI-only assignment.

Worker Welfare Analogously, we define worker (applicant) welfare as

$$W_A = \delta\,\mathbb{1}\{\hat{I}^m > \bar{I}\}, \qquad \delta > 0.$$

[6]Naturally ν and η can be made proportional to the gap between the applicants composite index from the cutoff.

That is, the worker receives payoff δ if she is employed—i.e., when her expected composite index $\hat{I}^m$ clears the hiring cutoff $\bar{I}$, and 0 otherwise.[7]

Overall Welfare Results As illustrated in Figure 4, when applicants can choose their screener, but the firm does not condition hiring decisions on this choice, three patterns emerge. First, firm welfare decreases with selection across all parameter values. Second, worker welfare increases under selection, across all parameter values. Third, the firm hires more high-quality applicants (fewer Type I errors) but also more low-quality applicants (more Type II errors).

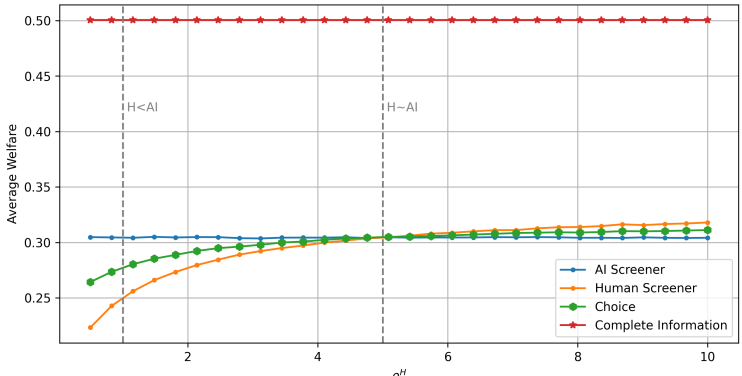
The intuition behind the firm's decline in welfare is as follows. Applicants naturally choose the screener that maximizes their chances of being hired. This may arise either because the chosen screener offers a higher expected assessment or because its variance is more favorable: high-ability applicants prefer lower-variance screeners, while low-ability applicants prefer higher-variance ones. This selection behavior distorts the distribution of signals that the firm observes. If the firm does not condition on which screener the applicant selected, it fails to account for this selection effect and, consequently, forms overly optimistic beliefs about candidates. The resulting misallocation leads to a clear reduction in firm welfare. Hence, if the firm cannot condition its decision on applicants' screener choices, it is strictly better off assigning a single screener rather than offering choice. As illustrated in Figure 4, firm welfare under choice always falls below the welfare that could be achieved by committing to the better of the two screeners.

The intuition behind the worker welfare results is as follows. When one screener dominates, high-type workers—those who satisfy the employment threshold $\alpha\theta_1 + (1-\alpha)\theta_2 > \bar{I}$—choose the more informative screener to reduce noise and better signal their true ability. In contrast, Low-type workers prefer the noisier screener, hoping that random variation will allow them to cross the cutoff. Consequently, both types benefit, resulting in higher overall employment. This pattern accounts for the observed decrease in Type I errors and the increase in Type II errors. A similar mechanism operates even when neither screener dominates. Workers select the screener that is more precise along the dimension in which they are strong, while relying on noise in the other dimension to obscure weaknesses. In both cases, being able to choose a screener increases the probability of exceeding the hiring threshold relative to being assigned a fixed screener.

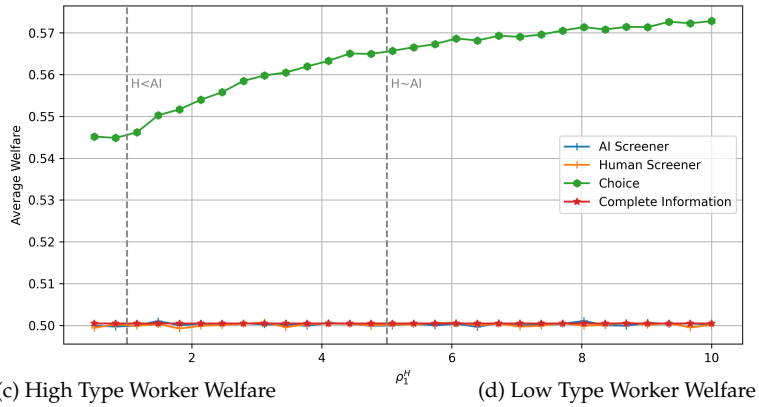
[7] Naturally, worker welfare could be extended to incorporate match quality, yielding higher utility when the worker truly clears the bar and lower utility otherwise. Here, however, we focus on the simpler case in which the worker's objective is solely to secure the job.

Figure 4: Welfare Under No Selection

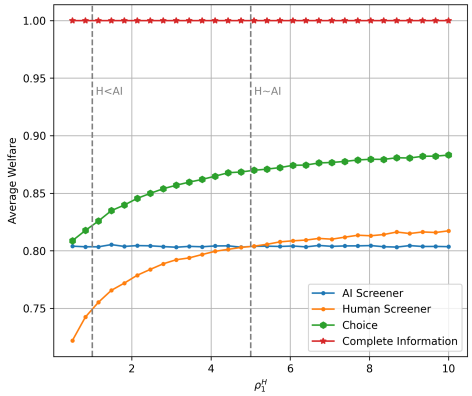
(a) Firm Welfare



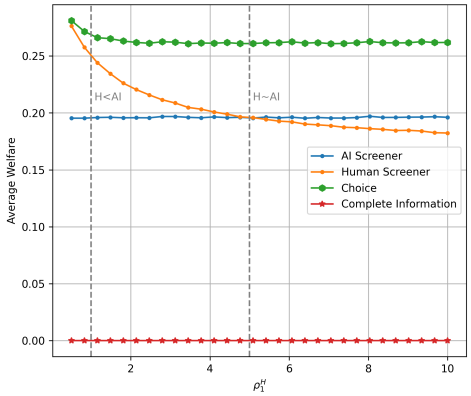
(b) Worker Welfare



(c) High Type Worker Welfare



(d) Low Type Worker Welfare



Notes: We generate $n = 1{,}000{,}000$ candidate draws for each ρ_1^H, with $\mu_1 = \mu_2 = 1$, $\rho_1 = \rho_2 = 1$, $\rho_2^H = 1$, $\rho_1^{AI} = 1$, $\rho_2^{AI} = 5$, $\bar{I} = 1$, and $\alpha = 0.5$, $\nu = \eta = \delta = 1$. The grid for ρ_1^H spans 30 values evenly spaced between 0.5 and 10. For values of ρ_1^H to the left of the vertical line labeled $H < AI$, the human screener is dominated by the AI screener. At the line labeled $H \sim AI$, the screeners have equal overall precision, though on opposite dimensions. We refer to workers whose attributes satisfy $\alpha\theta_1 + (1-\alpha)\theta_2 > \bar{I}$ as *high-type* workers, and to those below this threshold as *low-type* workers.

In sum, when the firm cannot condition on applicants' screener choices, it gains no additional information from offering a choice and thus experiences a welfare loss. In contrast, workers unambiguously benefit from having the ability to choose.

3.2 Firm Incorporates Selection

3.2.1 Firm's Inference and Hiring Decisions

Similarly to the case in which selection is ignored, after observing the signals from screener m the firm's posterior beliefs are

$$\theta_j \mid s_j^m \sim \mathcal{N}\left(\hat{\mu}_j^m, \hat{\rho}_j^m\right), \quad j \in \{1,2\}.$$

Where $\hat{\mu}_{ij}^m$ and $\hat{\rho}_j^m$ are as defined in Equation (1). However, the firm knows that the screener choice is driven by whether θ_1 is above or below $\beta_0 + \beta_1\theta_2$, where β_0 and β_1 are defined in Equation (4). Given the parameter values, upon observing the chosen screener, the firm knows whether the inequality holds. Assume $m = H$, write the constraint as

$$\theta_1 - \beta_1\,\theta_2 > \beta_0.$$

The random variable

$$L = \theta_1 - \beta_1\,\theta_2$$

is normally distributed with

$$\mathbb{E}[L] = \mathbb{E}[\theta_1|s_1^m] - \beta_1\,\mathbb{E}[\theta_2|s_2^m] = \frac{\rho_1\mu_1 + \rho_1^m s_1^m}{\rho_1 + \rho_1^m} \;-\; \beta_1\,\frac{\rho_2\mu_2 + \rho_2^m s_2^m}{\rho_2 + \rho_2^m},$$

$$\mathrm{Var}[L] = \mathrm{Var}[\theta_1|s_2^m] + \beta_1^2\,\mathrm{Var}[\theta_2|s_2^m] = \frac{1}{\rho_1 + \rho_1^m} \;+\; \frac{\beta_1^2}{\rho_2 + \rho_2^m}.$$

Let

$$\sigma_L = \sqrt{\frac{1}{\rho_1 + \rho_1^m} + \frac{\beta_1^2}{\rho_2 + \rho_2^m}}, \quad \psi = \frac{\beta_0 - \mathbb{E}[L]}{\sigma_L}.$$

The joint posterior density truncated to $\theta_1 - \beta_1\,\theta_2 > \beta_0$ is

$$f(\theta_1, \theta_2 \mid s_1^m, s_2^m,\, \theta_1 - \beta_1\theta_2 > \beta_0)$$

$$\propto\; \exp\!\Big[-\tfrac{\rho_1}{2}(\theta_1 - \mu_1)^2 - \tfrac{\rho_1^m}{2}(s_1^m - \theta_1)^2\Big] \;\times\; \exp\!\Big[-\tfrac{\rho_2}{2}(\theta_2 - \mu_2)^2 - \tfrac{\rho_2^m}{2}(s_2^m - \theta_2)^2\Big]$$

$$\text{subject to } \theta_1 - \beta_1\,\theta_2 > \beta_0,$$

normalized by the factor $\Phi(\psi)$, where Φ is the standard normal CDF. Below, we define the *inverse Mills ratio* for both the lower and upper tails of the standard normal distribution:

$$\lambda_-(\psi) \;=\; \frac{\phi(\psi)}{\Phi(\psi)} \quad \text{(lower–tail)}, \qquad \lambda_+(\psi) \;=\; \frac{\phi(\psi)}{1-\Phi(\psi)} \quad \text{(upper–tail)}.$$

This leads to the following expectations

$$\mathbb{E}[\theta_j \mid s_1^{m^*}, s_2^{m^*}, m^*] = \begin{cases} \hat{\mu}_j^{m^*} + \frac{\left(\mathbb{1}_{\{j=1\}} - \beta_1 \mathbb{1}_{\{j=2\}}\right)}{\rho_j + \rho_j^{m^*}} \frac{\left(\lambda_+(\psi)\mathbb{1}_{\{m^*=AI\}} - \lambda_-(\psi)\mathbb{1}_{\{m^*=H\}}\right)}{\sigma_L}, & \text{if } A_{AI} \geq A_H, \\ \hat{\mu}_j^{m^*} - \frac{\left(\mathbb{1}_{\{j=1\}} - \beta_1 \mathbb{1}_{\{j=2\}}\right)}{\rho_j + \rho_j^{m^*}} \frac{\left(\lambda_-(\psi)\mathbb{1}_{\{m^*=AI\}} - \lambda_+(\psi)\mathbb{1}_{\{m^*=H\}}\right)}{\sigma_L}, & \text{if } A_{AI} < A_H. \end{cases}$$

That is, in addition to forming expectations as before, $\hat{\mu}_j^{m^*}$, the firm now applies a correction that either inflates or deflates these expectations depending on the applicant's screener choice. When one screener dominates the other (i.e., when $\beta < 0$), an applicant who chooses the dominated screener causes the firm to downweight its expectations of both skill dimensions, whereas choosing the dominant screener leads the firm to raise its expectations across both dimensions.

When the screeners do not dominate each other (i.e., when $\beta > 0$), the correction becomes dimension-specific: depending on the applicant's choice, the firm increases its expectation for the dimension in which the chosen screener is more precise relative to the alternative and decreases the expectation for the other dimension.

In short, once the firm recognizes that applicants' choices are not random, it treats them as informative signals and incorporates this information into its belief-updating process.

Observation 2 (Firm's Beliefs Conditioning on Selection)**.**
(i) Unbiased Posterior Point Belief. *The firm's posterior mean beliefs, are unbiased point estimates for the true θ_j.*
(ii) Posterior Variance Reduction. *The Firm's posterior variance when m^* is endogenously chosen and accounted for is smaller than the variance if m^* had been exogenously assigned to the candidate.*

The supporting calculations are provided in Appendix Section A.1. Part (i) follows directly from Bayesian updating; however, unlike in the previous section—where the firm ignored the applicant's choice—its posterior beliefs are now centered at the true θ_j. For part (ii), the intuition is that the firm observes not only the signal generated by the chosen screener but also the applicant's choice itself, which restricts the set of possible types to a corresponding half-space. This additional informational constraint necessarily reduces the posterior variance relative to the case in which the screener is assigned exogenously.

3.2.2 Applicant's Problem

Since the applicant does not account for the firm internalizing the screener choice when forming posterior beliefs, the applicant's optimization problem remains unchanged from Section 3.1.2.

3.2.3 The Welfare Consequences of Conditioning Hiring on Screener Choice

As illustrated in Figure 5, when applicants can choose their screener *and* the firm conditions its hiring decision on this choice, several robust patterns emerge. First, firm welfare increases under selection, across all parameter values. Second, overall worker welfare depends on parameter values. Third, high-type workers benefit from selection across all parameter values. Fourth, low-type workers are harmed by selection across all parameter values. Fifth, the firm hires more high-quality applicants (fewer Type I errors) and fewer low-quality applicants (fewer Type II errors).

The intuition behind the firm's increase in welfare is as follows. Unlike in the previous setting, the firm now observes the applicant's screener choice and incorporates this information into its posterior beliefs. Even when one screener dominates the other, the firm gains informational value from seeing an applicant select the noisier screener: such a choice is itself a signal about the applicant's type. Thus, although offering choice may reduce the informativeness of the screener's signal, the firm gains an additional source of information through the selection decision. By combining the screening signal with this endogenous selection signal, the firm can make more accurate hiring decisions, thereby improving overall welfare. In particular, the firm hires more high-type applicants and fewer low-type applicants. As shown in Figure 5, firm welfare under endogenous selection exceeds what the firm could achieve by committing to either screener alone, and this advantage persists even as the human screener's precision ρ_1^H increases. While improvements in precision eventually generate diminishing returns, the informational content of the applicant's choice produces a persistent welfare gain.[8]

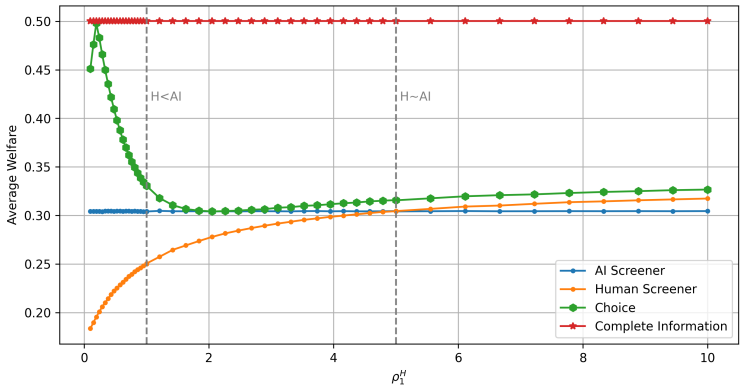
Whether workers overall gain or lose depends on parameter values. High-type workers benefit because the firm can identify them more effectively using both the signal and the selection decision, increasing their probability of being hired. By contrast, low-type

[8] Interestingly, in the region where the green (Choice) and red (Complete Information) curves nearly coincide, the cutoff boundary becomes highly informative. In this range, the estimated slope is negative, and the applicant's indifference line lies entirely below the firm's hiring threshold. Consequently, any applicant who selects the human screener effectively reveals that they fall below the hiring bar, enabling the firm to screen them out almost perfectly. This near-complete separation substantially increases efficiency, pushing firm welfare close to the full-information benchmark. Quantitatively, in this region, we find $\beta_0 \approx 2$ and $\beta_1 \approx -1$, indicating that the screener is nearly a perfectly revealing indicator of hire-worthiness.

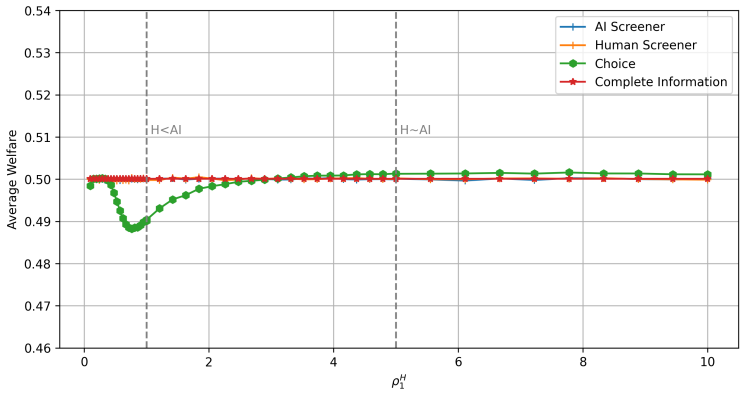
workers—those who do not meet the employment threshold—are worse off. With access to the selection signal, the firm becomes better not only at reducing Type I errors (hiring more high-type workers) but also at reducing Type II errors (hiring fewer low-type workers). Consequently, low-skilled workers succeed less often in obtaining a job and therefore experience a welfare loss.

Figure 5: Welfare Under Selection

(a) Firm Welfare

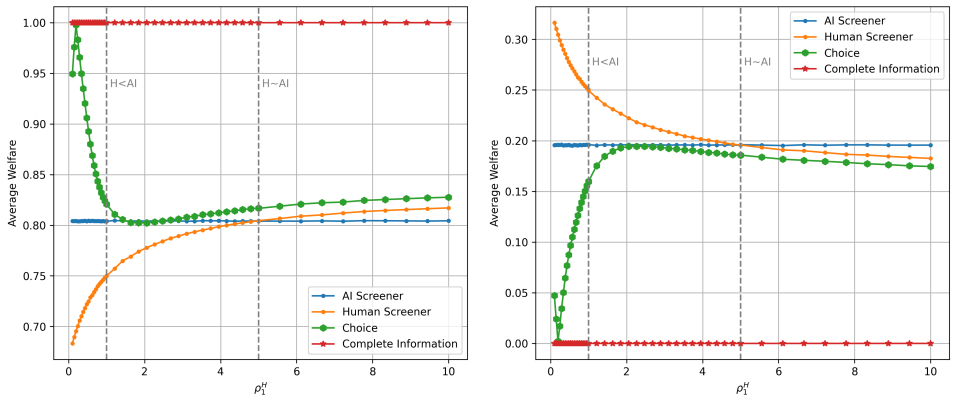


(b) Worker Welfare



(c) High Type Worker Welfare

(d) Low Type Worker Welfare



Notes: We generate $n = 10{,}000{,}000$ candidate draws for each ρ_1^H, with $\mu_1 = \mu_2 = 1$, $\rho_1 = \rho_2 = 1$, $\rho_2^H = 1$, $\rho_1^{AI} = 1$, $\rho_2^{AI} = 5$, $\bar{I} = 1$, and $\alpha = 0.5$, $\nu = \eta = \delta = 1$. The grid for ρ_1^H is constructed in three segments: 20 evenly spaced values from 0.1 to ρ_2^H, 20 values from ρ_2^H to ρ_2^{AI}, and 10 values from ρ_2^{AI} to $2\rho_2^{AI}$. For values of ρ_1^H to the left of the vertical line labeled $H < AI$, the human screener is dominated by the AI screener. At the line labeled $H \sim AI$, the screeners have equal overall precision, though on opposite dimensions. We refer to workers whose attributes satisfy $\alpha\theta_1 + (1-\alpha)\theta_2 > \bar{I}$ as *high-type* workers, and to those below this threshold as *low-type* workers.

In sum, when the firm can condition on applicants' screener choices, it gains additional information from offering a choice and thus achieves a welfare increase. High-type workers also benefit from the ability to choose, while low-type workers are adversely affected.

3.3 Mutual Incorporation of Selection

We next turn to the case in which both sides of the market incorporate the applicant's screener choice into their decision-making. In this regime, applicants anticipate that the firm will adjust its posterior beliefs based on their screener selection, and the firm, in turn, internalizes the informational content of this choice when forming expectations and setting hiring rules. The interaction between these two strategic responses has important implications for welfare and hiring outcomes.

3.3.1 Pooling Equilibria

In a pooling equilibrium all types (θ_1, θ_2) choose the same screener $m^* \in \{H, AI\}$. The applicant's strategy is therefore

$$\sigma(\theta_1, \theta_2) = m^* \quad \text{for all } (\theta_1, \theta_2) \in \mathbb{R}^2.$$

Since the choice of m^* conveys no additional information about type, the firm's poste-

rior mean of θ_j coincides with the update described in Section 3.1.1:

$$\mathbb{E}\left[\theta_j \,\middle|\, s_j^{m^*},\, m^*\right] = \frac{\rho_j\,\mu_j + \rho_j^{m^*}\, s_j^{m^*}}{\rho_j + \rho_j^{m^*}} = \hat{\mu}_j^{m^*}, \qquad j \in \{1,2\}.$$

The hiring decision is then based on the composite index

$$\hat{I}^{m^*} = \alpha\,\hat{\mu}_1^{m^*} + (1-\alpha)\,\hat{\mu}_2^{m^*},$$

with the applicant hired whenever $\hat{I}^{m^*} \geq \bar{I}$. For the unchosen screener $m' \neq m^*$, the firm must form off-path beliefs $\mathbb{E}\left[\theta_j \mid s_j^{m'}, m'\right]$ that prevent profitable deviations. A pooling equilibrium is sustained if, for every type (θ_1, θ_2),

$$\mathbb{P}\Big(\alpha\,\mathbb{E}[\theta_1 \mid s_1^{m'}, m'] + (1-\alpha)\,\mathbb{E}[\theta_2 \mid s_2^{m'}, m'] > \bar{I} \,\Big|\, \theta_1, \theta_2\Big) \;\leq\; \mathbb{P}\Big(\alpha\,\hat{\mu}_1^{m^*} + (1-\alpha)\,\hat{\mu}_2^{m^*} > \bar{I} \,\Big|\, \theta_1, \theta_2\Big),$$

with strict inequality for at least some types if the off-path probability is positive.

One simple supporting belief is that any applicant who selects m' is of arbitrarily low quality. Formally, by concentrating the firm's off-path prior on a set $\Theta_{m'}$ of very low types (e.g. $\theta_j \to -\infty$), the implied hiring probability collapses to zero,

$$\mathbb{P}(\text{Hired} \mid m') = 0,$$

which makes deviation to m' unprofitable.

Discussion of Pooling Equilibria While pooling equilibria are theoretically coherent and useful for understanding the model's possible outcomes, they are of limited practical relevance for our purposes. In the data, we observe sizable selection into both the human and AI screener, suggesting that applicants actively differentiate between options rather than collapsing into a single choice. Moreover, sustaining a pooling outcome requires somewhat contrived off-path beliefs, such as the assumption that any applicant who deviates to the other screener must be of substantially lower quality. For these reasons, while pooling equilibria help clarify the model's structure, our empirical focus and substantive interest lie in the richer separating equilibria that align with observed behavior.

3.3.2 Separating Equilibria

In a separating equilibrium, applicants of different types select different screeners, so that the type space is partitioned into subsets

$$\Theta_{AI}, \Theta_H \subseteq \mathbb{R}^2, \qquad \Theta_{AI} \cap \Theta_H = \emptyset,$$

with types in Θ_{AI} finding it optimal to choose the AI screener and those in Θ_H preferring the Human screener. Indifferences, when they occur, are assumed to be broken arbitrarily.

Upon observing $m \in \{H, AI\}$ and the realized signals s^m, the firm infers that the applicant belongs to the subset of types that would have selected that screener. The choice itself thus carries informational content, and the firm's posterior belief about ability reflects this selection event:

$$\mathbb{E}\left[\theta_j \,\middle|\, s^m, m\right] = \mathbb{E}\left[\theta_j \,\middle|\, s^m, \theta \in \Theta_m\right], \qquad j \in \{1, 2\}.$$

Hence, beliefs are updated using a version of the prior distribution truncated to the relevant subset of types. The hiring decision again follows a threshold rule based on the posterior composite index:

$$\alpha\, \mathbb{E}[\theta_1 \mid s^m, \theta \in \Theta_m] + (1 - \alpha)\, \mathbb{E}[\theta_2 \mid s^m, \theta \in \Theta_m] \;\geq\; \bar{I}.$$

Equilibrium conditions

(i) *Beliefs:* The firm's beliefs are Bayes-consistent given the sets (Θ_{AI}, Θ_H).

(ii) *Optimality:* Given the firm's rule, types in Θ_{AI} weakly prefer AI to H, and those in Θ_H weakly prefer H to AI.

(iii) *Fixed point:* The sets (Θ_{AI}, Θ_H) are self-consistent: given the firm's inference rule, applicants optimally choose precisely these sets. The firm's inferences and applicants' choices thus jointly form a stable equilibrium configuration.

While this formulation provides the most general characterization of a separating equilibrium, it is analytically intractable. Without restrictions on the geometry of the sets (Θ_{AI}, Θ_H), it is not even clear what form these sets might take, nor whether they have useful regularity properties. As a result, the firm's posterior beliefs can no longer be expressed in closed form, since the conditioning event $\theta \in \Theta_m$ does not admit a tractable representation. From a computational standpoint, simulating such an equilibrium is also

challenging: one would need to allow for arbitrarily shaped sets and determine how to numerically update beliefs, rather than using a known analytical formula.[9]

3.3.3 Separating Equilibrium with Fixed Adjustments

We now focus on an affine separating equilibrium, in which the type space is partitioned into two half-spaces. A convenient restriction since, as shown earlier, the firm's posterior beliefs are then available in closed form. There are two possible approaches to obtain such a structure. One is to derive it by imposing sufficient conditions under which applicants optimally adopt a half–space strategy in equilibrium. While we have identified conditions under which such a partition arises endogenously, we do not take that route. Those sufficient conditions are often cumbersome and not always stated in terms of primitives.[10] Moreover, the analytical and numerical computations required to verify these conditions, and to characterize the resulting equilibrium, are highly demanding for us as analysts and, consequently, unrealistic as a description of how firms or applicants actually form beliefs in practice within this environment.

Instead, we take a more pragmatic approach that preserves both tractability and realism. We assume that, in equilibrium, the firm incorporates the fact that applicants select their screener but does so through a simple adjustment: it applies a fixed *boost* or *penalty* to its posterior belief about each dimension θ_j based on the chosen screener m, independently of the observed signal. For example, upon seeing that an applicant chose the AI screener, the firm may increase its expectation of θ_1 while decreasing its expectation of θ_2, or vice versa, or shift both in the same direction—depending on the equilibrium thresholds used by applicants. Crucially, these adjustments do not depend on the realized signals, making the environment considerably simpler while still capturing the core intuition of how firms might incorporate selection in practice.

Firm's Inference and Hiring Decision Following the same steps as in Section 3.1 leads to the posterior mean, which we now adjust as follows

$$\tilde{\mu}_j^m(\theta_j) = \frac{\rho_j\,\mu_j + \rho_j^m\,\theta_j}{\rho_j + \rho_j^m} + \tau_j^m, \qquad (\tau_1^m, \tau_2^m) \in \arg\max_{(\tau_1,\tau_2)\in\mathbb{R}^2} W_f(\tau_1, \tau_2 | m).$$

[9] In principle, one could approximate the firm's inference using maximum-likelihood or similar estimation procedures, but doing so would blur the link between the theoretical model and the simulated equilibrium structure, making it difficult to interpret the results within the same framework.

[10] Even under the regularity properties afforded by Gaussian types and single–crossing arguments, the equilibrium characterization may depend on the curvature of the inverse Mills ratio or other higher–order terms, which are difficult to interpret economically and do not map cleanly to observable parameters.

Where the additional τ_j^m represents the adjustment the firm makes on dimension j after learning that the participant chose screener m, importantly, the firm chooses these adjustments to maximize its welfare, but chooses one such adjustment; thus, τ_j^m aims to correct for selection on average, rather than for each observed signal. In total there will be four adjustments $(\tau_1^H, \tau_2^H, \tau_1^{AI}, \tau_2^{AI})$ either one allowed to be positive or negative.

Applicant's Problem Because in equilibrium τ_j^m will be a constant, it only shifts the mean of the composite index, which continues to follow a normal distribution

$$\hat{I}^m = \alpha\,\hat{\mu}_1^m + (1-\alpha)\,\hat{\mu}_2^m \sim \mathcal{N}\Big(\alpha\,\tilde{\mu}_1^m + (1-\alpha)\,\tilde{\mu}_2^m,\ \alpha^2\,\tilde{\rho}_1^m + (1-\alpha)^2\,\tilde{\rho}_2^m\Big).$$

The only difference from Section 3.1 being the addition of τ_j^m within $\tilde{\mu}_j^m(\theta_j)$. The applicant thus once again calculates

$$p_m(\theta) = \Pr\big(\hat{I}^m \geq \bar{I} \mid \theta\big) = 1 - \Phi\left(\frac{\bar{I} - \big(\alpha\,\tilde{\mu}_1^m + (1-\alpha)\,\tilde{\mu}_2^m\big)}{\sqrt{\alpha^2\,\tilde{\rho}_1^m + (1-\alpha)^2\,\tilde{\rho}_2^m}}\right).$$

where $p_m(\theta)$ is the probability of hiring under the screening m given type $\theta = (\theta_1, \theta_2)$. Recall

$$p_m(\theta) = \Pr\big(\hat{I}^m \geq \bar{I} \mid \theta\big) = 1 - \Phi\left(\frac{\bar{I} - \left(\alpha\left(\frac{\rho_1\,\mu_1 + \rho_1^m\,\theta_1}{\rho_1 + \rho_1^m} + \tau_1^m\right) + (1-\alpha)\left(\frac{\rho_2\,\mu_2 + \rho_2^m\,\theta_2}{\rho_2 + \rho_2^m} + \tau_2^m\right)\right)}{\sqrt{\alpha^2\,\frac{\rho_1^m}{(\rho_1+\rho_1^m)^2} + (1-\alpha)^2\,\frac{\rho_2^m}{(\rho_2+\rho_2^m)^2}}}\right).$$

Once again, since $\Phi(\cdot)$ is strictly increasing, the applicant simply compares the arguments of the normal CDF across screeners and chooses the one with the lower value. To characterize the indifference condition, for $m \in \{AI, H\}$ we utilize R_m, A_m, and B_m as defined in Equation (3). Note that the parameters τ_1^m and τ_2^m enter both the firm's and applicants' decisions only through the weighted combination $\alpha\tau_1^m + (1-\alpha)\tau_2^m$. Hence, we can define a composite adjustment term

$$\tau^m \equiv \alpha\tau_1^m + (1-\alpha)\tau_2^m,$$

Thereby reducing the problem to one in which the firm effectively chooses a single adjustment parameter for each screener.

$$\tau^m \in \arg\max_{(\tau)\in\mathbb{R}} W_f(\tau_1, \tau_2 | m)$$

The applicant is indifferent between AI and H along a linear cutoff in (θ_1, θ_2)-space:

$$\theta_1 = \beta_0 + \beta_1 \theta_2,$$

with the new intercept coefficient

$$\tilde{\beta}_0 \;=\; \frac{\bar{I}\,(R_{AI}^{-1} - R_H^{-1}) - \left(\frac{\tau_{AI}}{R_{AI}} - \frac{\tau_H}{R_H}\right) - \rho_1\mu_1\left(\frac{A_{AI}}{\rho_1^{AI}} - \frac{A_H}{\rho_1^{H}}\right) + \rho_2\mu_2\left(\frac{B_H}{\rho_2^{H}} - \frac{B_{AI}}{\rho_2^{AI}}\right)}{A_{AI} - A_H},$$

While the slope remains unchanged from

$$\beta_1 \;=\; \frac{B_H - B_{AI}}{A_{AI} - A_H}.$$

Thus, the optimal screener choice, once again, is determined by which side of this cutoff line the applicant's type lies:

$$m^*(\theta) = \begin{cases} \text{AI}, & \text{if } \left(A_{AI} \geq A_H \;\text{ and }\; \theta_1 \geq \tilde{\beta}_0 + \beta_1\theta_2\right) \;\text{ or }\; \left(A_{AI} < A_H \;\text{ and }\; \theta_1 \leq \tilde{\beta}_0 + \beta_1\theta_2\right), \\ H, & \text{otherwise.} \end{cases}$$

Equilibrium conditions

(i) *Beliefs:* The firm forms Bayes–consistent posterior beliefs adjusted by fixed parameters (τ^H, τ^{AI}), which represent its average corrections for self–selection under each screener.

(ii) *Optimality:* Given these adjustments, applicants optimally choose between H and AI according to the induced half–space rule, selecting the screener that yields a higher hiring probability.

(iii) *Fixed point:* The adjustments (τ^H, τ^{AI}) and the resulting cutoff are jointly self–consistent: given the firm's inference rule, applicants' choices reproduce exactly the separation that justifies those same adjustments. The firm's inference and applicants' responses thus constitute an affine separating equilibrium with fixed adjustments.

We describe the simulation procedure in Appendix Section B.

3.3.4 The Welfare Consequences of Mutual Incorporation of Selection

As shown in Figure 6, when applicants can choose their preferred screener *and* it is common knowledge that the firm will condition hiring decisions on this choice, four equi-

librium patterns arise. First, the firm weakly gains from allowing applicants to select a screener. Second, high-type workers benefit only weakly from selection. Third, low-type workers are weakly harmed by selection. Fourth, the firm weakly hires more high-quality applicants (fewer Type I errors) and fewer low-quality applicants (fewer Type II errors).

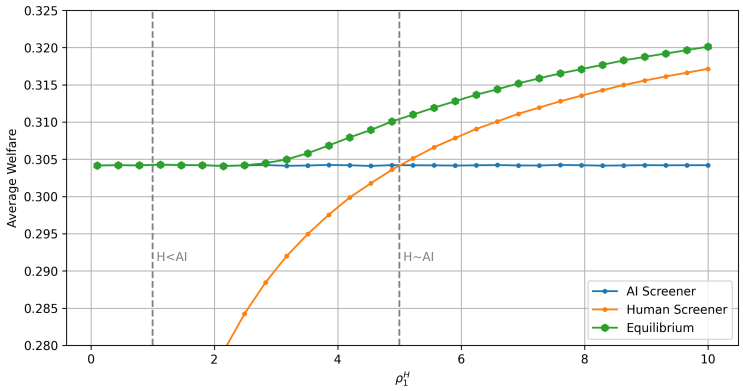
The intuition behind these results is the following. When one screener is substantially less precise than the other—e.g., when ρ_1^H is below 3 in Figure 6—choosing the human screener effectively signals that the candidate may be relying on noise to secure employment. Anticipating this, the firm significantly downgrades its expectations of such applicants. Knowing this, candidates avoid selecting the low-precision screener altogether. Thus, if one screener is clearly inferior, it is never chosen in equilibrium, and both firm and worker welfare revert to the benchmark in which only the more precise screener exists. In the figure, this appears as the green and blue curves coinciding for low ρ_1^H.

In contrast, when both screeners are reasonably precise and neither dominates the other, candidates' choices reflect their comparative advantage across skill dimensions rather than reliance on noise. High-θ_1 applicants may choose the screener that is more precise in assessing θ_1, and similarly for θ_2. Observing this, the firm updates its beliefs not only on the basis of direct signals but also on the basis of the choice itself. This additional information improves sorting: the firm becomes more likely to hire high-type workers and less likely to hire low-type workers, thereby reducing both Type I and Type II errors and raising firm welfare.

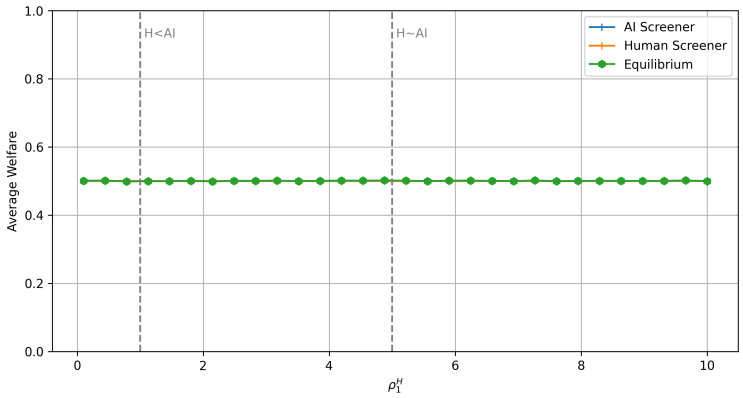
For workers, high-type applicants are more likely to be hired, while low-type applicants are less likely to be hired. On average, these gains and losses offset each other, leaving total worker welfare roughly unchanged, although the distributional consequences remain meaningful.

Figure 6: Equilibrium Welfare

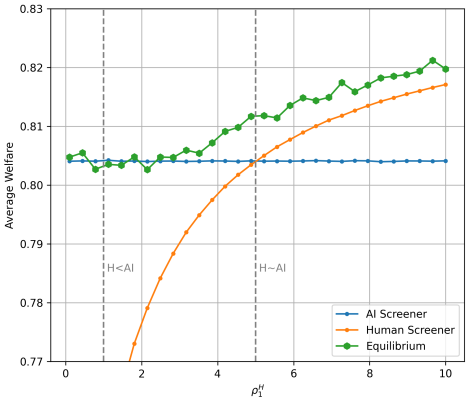
(a) Firm Welfare



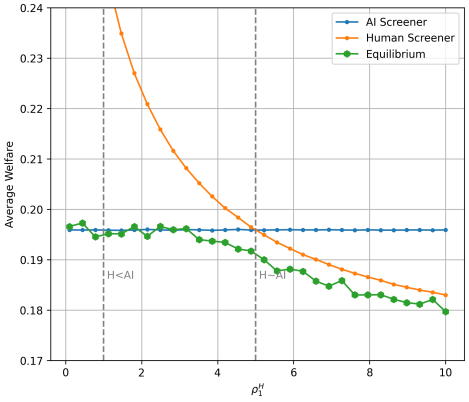
(b) Worker Welfare



(c) High Type Worker Welfare



(d) Low Type Worker Welfare



Notes: We generate $n = 10{,}000{,}000$ candidate draws for each ρ_1^H, with $\mu_1 = \mu_2 = 1$, $\rho_1 = \rho_2 = 1$, $\rho_2^H = 1$, $\rho_1^{AI} = 1$, $\rho_2^{AI} = 5$, $\bar{I} = 1$, and $\alpha = 0.5$, $\nu = \eta = \delta = 1$. The grid for ρ_1^H spans 30 values evenly spaced between 0.5 and 10. For each case, we find the equilibrium following the simulation procedure described in Appendix Section B. For values of ρ_1^H to the left of the vertical line labeled $H < AI$, the human screener is dominated by the AI screener. At the line labeled $H \sim AI$, the screeners have equal overall precision, though on opposite dimensions. We refer to workers whose attributes satisfy $\alpha\theta_1 + (1-\alpha)\theta_2 > \bar{I}$ as *high-type* workers, and to those below this threshold as *low-type* workers.

In summary, when the firm conditions on screener choice and this is common knowledge, an inferior screener goes unused in equilibrium, and welfare mirrors the single-screener case. When, instead, both screeners specialize along different dimensions, both are used in equilibrium, the firm and high-ability workers benefit, and low-ability workers are hired less often.

Choice as a policy regulation for AI adoption in the labor market. Our design also addresses current debates on how regulators should govern AI adoption in the labor market. One natural proposal is to grant applicants a *right to choose* whether they are screened by a human or by an AI system, echoing emerging 'right to human review" provisions in data-protection and AI regulation.

Our framework makes clear that such a right can be implemented in multiple ways. In one regime, applicants can choose their preferred interviewer, but firms are prohibited from conditioning hiring decisions on that choice; this maximizes decision autonomy but deliberately discards the informational content of selection into human versus AI screening. In a second regime, firms are permitted to treat interviewer choice as an additional signal about applicant type; this may improve match quality and firm welfare, but it also creates systematic differences across applicants depending on how their choices are inferred.[11] Our estimates quantify these trade-offs in a large-scale, real-world screening environment.

While we do not attempt to solve the regulator's optimal policy problem, our results show that offering choice is not a purely procedural add-on. Depending on how firms are allowed to use the resulting choice signal, it can materially shift the distribution of welfare between firms and applicants and across applicant types.

4 Estimation and Calibration Procedure

In this section, we treat the language-specific component of worker ability as the first dimension, θ_1, and the analytical component as the second, θ_2. Figure 7 summarizes our identification and estimation strategy. Each block corresponds to a set of moments in the data and the subset of parameters that they identify. We begin by describing the initial procedure (Steps 1–5), which uses data from the deterministic human-AI systems to obtain the preliminary parameter estimates reported in Table 2. We then present the comprehensive estimation procedure in Section 4.3.

One difference between this section and the theory section is that, for each candidate, the firm also receives scores from a language test and an analytical test. These can be readily incorporated into our framework as additional signals about the corresponding dimensions. We then incorporate these signals into the firm's posterior beliefs and, in turn, into its hiring decisions, retention rates, and related outcomes. Our framework is flexible and can be easily adapted to accommodate industry-specific characteristics such as this one.

[11] The practical feasibility of these regimes will depend on the institutional setting. If firms can directly observe an applicant's choice of interviewer, no regulation can fully prevent them from incorporating that information into their inference. In practice, implementing a regime that restricts the use of choice as a signal may require an intermediary who evaluates interview outcomes without being informed of the applicant's selection. Designing workable institutional arrangements along these lines is an important task for future research.

Data Used in Estimation Our estimation uses the randomized sample described in Section 2.3. For the initial procedure in Section 4.1, we restrict to applications assigned to the deterministic Human-AI systems and focus on the following variables from the recruitment system described in Section 2.2: (i) the interviewer system $m \in \{\text{H}, \text{AI}\}$ and interview score IS^m, described there and rescaled to have mean zero and unit variance within system; (ii) the standardized language and analytical test scores (LS, AS) defined in the "Standardized language and analytical tests" paragraph; (iii) the job-offer indicator from the review stage; and (iv) retention at horizon h months as defined in the "Downstream employment outcomes" paragraph. Other features of the environment (engagement-based pre-screening, scheduling, and on-the-job productivity metrics) are important for external validity and for our reduced-form evidence, but they do not enter the structural moments we construct in this section.

4.1 Initial Estimation Procedure

In steps (1-5) that follow, we use data from the deterministic human and deterministic AI treatments, where applicants were directly assigned to a screener.

Step 1: Normalization We impose the normalization $\mu_1 = \mu_2 = 0$ and $\rho_1 = \rho_2 = 1$. Normalizing the means to zero entails no loss of generality. Allowing μ_1 and μ_2 to differ would prevent point identification of the parameter α, as any mean differences would be absorbed by it. Similarly, normalizing $\rho_1 = \rho_2 = 1$ does not restrict the generality of the model, since changes in these parameters would shift the estimated precision of the interview screeners, ρ_j^m, proportionally. This corresponds to the *Normalization* block in Figure 7, which fixes the scale and location of the latent ability distribution (θ_1, θ_2) and the prior precisions (ρ_1, ρ_2) that define the units in which all subsequent parameters are identified.

Step 2: Estimating Test Precision and Interview Score Weights In addition to information generated through the interviews, the firm observes noisy signals of the two ability dimensions through a language test score (LS) and an analytical test score (AS). As detailed in Section 2.2, the language module maps performance into CEFR bands (A1–C2), and the analytical module aggregates three timed components—attention to detail, verbal reasoning, and numerical ability into a 0–100 composite score; we transform these raw scores into standardized measures LS and AS, which we model as unbiased Gaussian sig-

nals about θ_1 and θ_2 with common precision κ:[12]

$$LS \coloneqq t_1 \sim \mathcal{N}\left(\theta_1, \frac{1}{\kappa}\right), \qquad AS \coloneqq t_2 \sim \mathcal{N}\left(\theta_2, \frac{1}{\kappa}\right).$$

In addition to the test scores, the interviews also yield informative signals regarding both dimensions, which we have denoted as s_1^m and s_2^m. We model the interview score (IS^m) as a linear combination of these two underlying signals:

$$IS^m = \delta_1^m s_1^m + \delta_2^m s_2^m$$

Running a regression of IS^m on t_1 and t_2 yields estimated parameters

$$\hat{\beta}_{LS,m} = \frac{\text{Cov}(t_1, IS^m)}{\text{Var}(t_1)} = \frac{\delta_1^m}{1+\frac{1}{\kappa}}, \qquad \hat{\beta}_{AS,m} = \frac{\text{Cov}(t_2, IS^m)}{\text{Var}(t_2)} = \frac{\delta_2^m}{1+\frac{1}{\kappa}}.$$

Which give us an equation linking the test precision with the weights used on each interview signal. This corresponds to the *Regression Analysis Linking IS, LS, and AS* block in Figure 7, which provides moments linking the test precision κ to the interview-score weights (δ_1^m, δ_2^m) and hence to the relative informativeness of the two underlying interview signals.

Step 3: Utilizing Hiring Decisions Conditional on Scores Given s_1^m, s_2^m, t_1, and t_2, the firm's posterior expectation of the worker's composite ability is

$$\begin{aligned}
\mathbb{E}[I \mid IS^m, t_1, t_2] &= \alpha\left(\frac{\kappa t_1 + \rho_1^m s_1^m}{1+\kappa+\rho_1^m}\right) + (1-\alpha)\left(\frac{\kappa t_2 + \rho_2^m s_2^m}{1+\kappa+\rho_2^m}\right) \\
&= \alpha\left(\frac{\rho_1^m}{1+\kappa+\rho_1^m}\right) s_1^m + (1-\alpha)\left(\frac{\rho_2^m}{1+\kappa+\rho_2^m}\right) s_2^m \\
&\quad + \alpha\left(\frac{\kappa}{1+\kappa+\rho_1^m}\right) t_1 + (1-\alpha)\left(\frac{\kappa}{1+\kappa+\rho_2^m}\right) t_2.
\end{aligned}$$

Under the assumption that the interview score is constructed in accordance with the firm's posterior weights on s_1^m and s_2^m, we have

$$\delta_1^m = \alpha\left(\frac{\rho_1^m}{1+\kappa+\rho_1^m}\right), \qquad \delta_2^m = (1-\alpha)\left(\frac{\rho_2^m}{1+\kappa+\rho_2^m}\right).$$

[12] This assumption can be relaxed in future work, for example by adopting a nonparametric specification of the test score distributions, which we directly observe. In that case, it is also straightforward to allow each test to have its own precision parameter, as these would be disciplined by the observed distributions.

That is, when forming the interview score, each dimension is weighted proportionally to its relative precision and to the importance the firm assigns to that dimension.[13]

We next estimate a probit model of the conditional probability of hiring a candidate given IS^m, t_1, and t_2. The hire indicator in this probit regression corresponds to the recruiter decision described inSection 2.2, where a human reviewer observes both the interview and test results and issues a job offer. The estimated coefficients on t_1 and t_2, the parameters of interest that help us discipline our primitives, will then be

$$\hat{\beta}^{prob}_{t_1,m} = \alpha \frac{\kappa}{1+\rho_1^m+\kappa}, \qquad \hat{\beta}^{prob}_{t_2,m} = (1-\alpha)\frac{\kappa}{1+\rho_2^m+\kappa}.$$

However, if we do not assume that the errors in the probit model are standard normal, the coefficients would be scaled by an additional term. To eliminate this dependence, we can take the ratio of the two parameters, which cancels out any normalization of the latent error variance. This ratio therefore provides an additional moment condition linking κ, ρ_1^m, ρ_2^m, and α.:

$$\frac{\hat{\beta}^{prob}_{LS}}{\hat{\beta}^{prob}_{AS}} = \frac{\alpha}{1-\alpha}\frac{1+\rho_2^m+\kappa}{1+\rho_1^m+\kappa}.$$

This corresponds to *the conditional job offer score* block in Figure 7, which uses the variation in the hiring decisions conditional on (IS^m, t_1, t_2) to identify α, κ, and the relative precision of the language and analytical dimensions for each screener, ρ_1^m/ρ_2^m.

We use the steps described above to identify κ and the relative weight of ρ_1^m compared to ρ_2^m. We then proceed with the steps below to identify $\bar{I}$ and α, and rescale ρ_1^m and ρ_2^m by a common constant to match the data, while preserving their previously established relative signal precisions.

Step 4: Unconditional Hiring Probability We next exploit unconditional job offer data to provide additional restrictions linking the parameters of interest. Under Assumption 1, 2, and 3, it follows that as the precision of the screeners increases, the probability of employment also rises, which is reflected in a higher unconditional offer rate.[14] Formally,

[13]Given the available data, this restriction can be relaxed in subsequent analysis. For a particular example, if $\delta_1^m = \gamma^m$, and $\delta_2^m = 1-\gamma^m$, then the regression under **Step 2** directly point identifies both κ and γ^m.

[14]To see this note that $\mathbb{E}[\hat{\mu}_j] = \mu_j$ and $\text{Var}[\hat{\mu}_j] = \frac{\kappa+\rho_j^m}{(\rho_j+\kappa+\rho_j^m)\rho_j}$. Then, under the assumption that $\bar{I} > \alpha\mu_1 + (1-\alpha)\mu_2$, the argument of $\Phi(\cdot)$ is strictly decreasing in each ρ_j^m, implying that the probability of hire is strictly increasing in the screeners' precision.

for a given screening technology $m \in \{H, AI\}$, the unconditional probability of hire is

$$P(\text{Hire} \mid m) = 1 - \Phi\left(\frac{\bar{I} - \alpha\mu_1 - (1-\alpha)\mu_2}{\sqrt{\alpha^2 \frac{\kappa+\rho_1^m}{(\rho_1+\kappa+\rho_1^m)\rho_1} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{(\rho_2+\kappa+\rho_2^m)\rho_2}}}\right)$$
$$= 1 - \Phi\left(\frac{\bar{I}}{\sqrt{\alpha^2 \frac{\kappa+\rho_1^m}{1+\kappa+\rho_1^m} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{1+\kappa+\rho_2^m}}}\right).$$

The second equality follows from our normalization. Empirically, we can replace the left-hand side with observed unconditional offer rates under human ($m = H$) and AI-based ($m = AI$) screening, so that $P(\text{Hire} \mid m)$ is the fraction of applications in systems m for which a job offer is issued. Since $\bar{I}$, α, and κ are invariant across screening modes, cross-screener differences in observed hiring rates directly reflect differences in ρ_j^m, the relative informativeness of human and AI screening signals. This corresponds to the *Unconditional Job Offers* block in Figure 7, which matches the overall offer rates under human and AI screening to determine the common hiring threshold $\bar{I}$, the weights the Firm places on each dimension α, and the screener precision levels $(\rho_1^H, \rho_2^H, \rho_1^{AI}, \rho_2^{AI})$, consistent with the relative magnitudes already identified in steps above.

Step 5: Unconditional Retention Data Next we use retention data from the deterministic human and the deterministic AI treatments as additional moments restricting parameters of interest. We measure $P(\text{Retention} \mid m)$ using the de-identified employment records described in Section 2.2, focusing on retention at 3 months as our baseline horizon. In the high-turnover BPO environment documented in Section 2—with annual separation rates around 50–60%—retention at these horizons is a natural proxy for worker–firm match quality, consistent with the labor literature cited there. Let

$$T = \alpha\theta_1 + (1-\alpha)\theta_2, \qquad B = \alpha\hat{\mu}_1 + (1-\alpha)\hat{\mu}_2.$$

T denotes the true composite index the firm would form under complete information, while B represents the posterior beliefs, or the composite index formed under imperfect information and observed signals. Their variances are

$$\sigma_T^2 = \alpha^2 + (1-\alpha)^2, \qquad \sigma_B^2 = \alpha^2 \frac{\kappa+\rho_1^m}{1+\kappa+\rho_1^m} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{1+\kappa+\rho_2^m}.$$

Because $B = \mathbb{E}[T \mid s_1^m, s_2^m, t_1, t_2]$, the covariance satisfies

$$\mathrm{Cov}(T, B) = \sigma_B^2,$$

and therefore the correlation between T and B is

$$\mathrm{Corr}(T, B) = \frac{\mathrm{Cov}(T, B)}{\sigma_T \sigma_B} = \frac{\sigma_B}{\sigma_T}.$$

Therefore,

$$\mathbb{P}(T > I \mid B > I) = \frac{1 - \bar{\Phi}_2\left(\frac{I}{\sigma_T}, \frac{I}{\sigma_B}; \frac{\sigma_B}{\sigma_T}\right)}{1 - \bar{\Phi}\left(\frac{I}{\sigma_B}\right)},$$

where $\bar{\Phi}$ is the standard normal cdf and $\bar{\Phi}_2(\cdot, \cdot; r)$ is the bivariate standard normal cdf with correlation σ_B/σ_T. So we further pin down the parameters via

$$\begin{aligned} P(Retention|m) &= \mathbb{P}(\alpha\theta_1 + (1-\alpha)\theta_2 > \bar{I} | \alpha\hat{\mu} + (1-\alpha)\hat{\mu}_2 > \bar{I}) \\ &= \frac{1 - \bar{\Phi}_2\left(\frac{\bar{I}}{\sqrt{\alpha^2 + (1-\alpha)^2}}, \frac{\bar{I}}{\sqrt{\alpha^2 \frac{\kappa+\rho_1^m}{(1+\kappa+\rho_1^m)} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{(1+\kappa+\rho_2^m)}}}; \frac{\sqrt{\alpha^2 \frac{\kappa+\rho_1^m}{1+\kappa+\rho_1^m} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{1+\kappa+\rho_2^m}}}{\sqrt{\alpha^2 + (1-\alpha)^2}}\right)}{1 - \bar{\Phi}\left(\frac{\bar{I}}{\sqrt{\alpha^2 \frac{\kappa+\rho_1^m}{(1+\kappa+\rho_1^m)} + (1-\alpha)^2 \frac{\kappa+\rho_2^m}{(1+\kappa+\rho_2^m)}}}\right)}. \end{aligned}$$

This corresponds to the *Unconditional Retention Rates* block in Figure 7, which uses differences in retention across human-AI systems to discipline the mapping from posterior beliefs B to true performance T and further refine the estimates of $\bar{I}$, α, and the screener precision parameters (ρ_1^m, ρ_2^m). When forming the moment above, we assume that retention is driven by match quality: the firm retains a worker if they are truly a good match and dismisses them otherwise. Because retention is measured at least one month after hiring, we implicitly assume that by that time the firm has effectively learned the worker's true type—the same latent attributes it attempted to infer at the hiring stage. Accordingly, if the firm learns that a worker's true θ_1 and θ_2 fall below the level it seeks, it lays the worker off.

There is, however, an important nuance. Some separations in the data arise because workers voluntarily leave the firm rather than being dismissed. Since the firm does not fire these workers, a natural interpretation is that their ability is high enough to justify

their initial hiring.[15] In the main analysis below, we introduce a parameter v that captures the share of voluntarily departing workers we classify as high-ability (i.e., workers who would have exceeded the hiring composite index under perfect information). In the baseline estimation reported in the main text, we set $v = 1$, consistent with the arguments above. As a robustness exercise, the appendix presents preliminary estimations with $v = 0.75$ and $v = 0.5$.

Summary of Moments Table 1 summarizes how moments from the deterministic human and deterministic AI systems map into the parameters identified in the initial procedure, mirroring the blocks in Figure 7.

Table 1: Summary of Moments Used in the Initial Procedure

Moments	Parameters primarily disciplined
Distribution of (LS, AS) and IS^m	κ, (δ_1^m, δ_2^m)
Score-conditional hire probs (IS^m, LS, AS)	κ, ratios ρ_1^m / ρ_2^m
Unconditional hire rates by m	$\bar{I}$, α, levels (ρ_1^m, ρ_2^m)
Unconditional retention rates by m	$\bar{I}$, α, levels (ρ_1^m, ρ_2^m)

Preliminary Estimated Parameters Following the above procedure, our estimates of the models parameters are as follows.

Table 2: Preliminary Estimated Parameters

$\hat{\bar{I}}$	$\hat{\alpha}$	$\hat{\rho}_1^H$	$\hat{\rho}_2^H$	$\hat{\rho}_1^{AI}$	$\hat{\rho}_2^{AI}$	$\hat{\kappa}$
1.12	0.88	7.95	7.55	7.86	9.66	0.32

Notes: The estimated parameters above result from following the steps outlined in Section 4.1, which combine estimation and calibration exercises. These parameters are subject to revision.

The relevant regressions are provided in Appendix Section C.1, and alternative estimations in Appendix Section C.2.

Interpretation of Preliminary Estimates Table 2 reveals several qualitative patterns. First, the estimated hiring threshold $\hat{\bar{I}} = 1.12$ implies a relatively selective rule: under our normalization, this corresponds to hiring workers whose expected composite ability

[15] Moreover, voluntary departures may be interpreted as a signal that the worker has secured a better outside option, further reinforcing the idea that they have relatively high ability.

is roughly 1.12 standard deviations above the population mean. Second, the estimated weight on the language dimension is $\hat{\alpha} = 0.88$, indicating that the composite index places substantial emphasis on language ability relative to analytical ability. Third, when comparing human and AI screeners, the human interviewer appears slightly more precise on the language dimension ($\hat{\rho}_1^H > \hat{\rho}_1^{AI}$), whereas the AI screener is more precise on the analytical dimension ($\hat{\rho}_2^{AI} > \hat{\rho}_2^H$).[16]

Finally, the precision parameters indicate that interview-based signals are far more informative than the standardized tests. There are different ways to interpret this. One possibility is that the interviews provide much more precise measures of these underlying dimensions. Alternatively, it may be that the firm is not primarily responding to the specific language and analytical skills captured by the standardized tests. Rather, there are adjacent competencies—such as communication clarity, situational judgment, or the ability to articulate ideas—that the interview captures more effectively than a basic language test. Similarly, it may not be the exact computations required in the analytical test that the firm values, but broader abilities such as logical reasoning or problem-solving, for which interviews may offer a more accurate read. In either case, our preliminary estimations suggest that most of the screening power comes from the interviews rather than the tests.

We view this exercise as a proof of concept intended to demonstrate that our model can be estimated using the available data. The specific coefficient estimates should not be taken at face value; we caution that these are preliminary results derived from an early-stage estimation and calibration exercise, and are likely to change as we carry out the full estimation procedure laid out next in Section 4.3.

Implications for Hybrid Screeners The retention mapping in Step 5 implies that more informative screening—a higher correlation between the posterior index B and true match quality T—strictly raises retention among hired workers, $P(T > \bar{I} \mid B > \bar{I})$. Given the preliminary estimates in Table 2, both human and AI interviews are highly informative relative to the standardized tests, and each screener is comparatively stronger on a different dimension of ability. A system that observes *both* human and AI interview signals therefore attains a higher effective precision on each dimension than either stand-alone system, which increases σ_B and hence the correlation $\text{Corr}(T, B) = \sigma_B/\sigma_T$. Under our assumptions, the model thus predicts that combining human and AI screeners should yield higher retention among hires than relying on either screener alone. The preliminary counterfactual exercises in Section 4.2 confirm this prediction numerically.

[16]This pattern motivates the dimension-specific counterfactuals below, in which each dimension is assigned to the screener that appears to measure it most accurately.

Next Steps in Estimation The approach described above combines estimation and calibration, providing a proof of concept that our model is identifiable from the available data. As such, we view the above estimated parameters as preliminary results subject to change. As a next step, we plan to estimate the model parameters using a joint maximum-likelihood procedure that incorporates both hiring and retention data, conditional on candidates' test and interview scores. This approach will allow us to estimate all parameters simultaneously and compute their standard errors.

4.2 Preliminary Counterfactuals

In this subsection, we use the preliminary parameter estimates from Section 4.1 to study simple counterfactual changes in interviewer assignment within the deterministic environment. We focus on two designs: a hybrid system in which each applicant is evaluated by both a human and an AI screener, and a dimension-specific system in which humans and AI are assigned to the dimensions where they are most precise. Holding the hiring threshold $\bar{I}$ fixed, these counterfactuals isolate how changing the information structure—while keeping selectiveness constant—affects job offers and retentions, as measured by involuntary separations. We deliberately abstract here from the Choice system and order effects; counterfactuals that explicitly involve choice and interview sequencing are developed after the comprehensive estimation in Section 4.3.

4.2.1 Hybrid Human *and* AI Screener

With the preliminary parameter estimates in hand, we consider a system in which each candidate is evaluated by both the human *and* the AI screener, in sequence, and the firm observes all corresponding interview signals.[17] For this exercise, the firm updates its beliefs using signals from language and test scores with precision κ, signals from the human screener with precision ρ_j^H, and signals from the AI screener with precision ρ_j^{AI}. The firm therefore forms expectations of the composite index as follows:

$$\mathbb{E}[I \mid s_1^H, s_2^H, s_1^{AI}, s_2^{AI}, t_1, t_2] = \alpha \left(\frac{\kappa t_1 + \rho_1^H s_1^H + \rho_1^{AI} s_1^{AI}}{1 + \kappa + \rho_1^H + \rho_1^{AI}} \right) + (1 - \alpha) \left(\frac{\kappa t_2 + \rho_2^H s_2^H + \rho_2^{AI}}{1 + \kappa + \rho_2^H + \rho_2^{AI}} \right),$$

[17] We do not implement such a double-interview system in our field experiment, but it is a natural design variation the partner firm could adopt in future work, randomizing whether the human or the AI screener sees the candidate first. Within our framework, this corresponds to an additional counterfactual in which the effective precisions (ρ_j^H, ρ_j^{AI}) and/or the hiring threshold $\bar{I}$ are allowed to depend on the interview order. Relatedly, Radbruch and Schiprowski (2025) study how the sequence in which human recruiters interview candidates affects subjective assessments in high-stakes admission and hiring processes.

and the firm employs the worker whenever this posterior exceeds the threshold $\bar{I}$. Unconditional job offers and retention rates in this hybrid system can then be computed using the same formulas as in the deterministic human-AI systems, replacing the posterior composite index used there with the one above.

To quantify the gains from combining the human and AI screeners, we compare the hybrid screening approach against each stand-alone screener. The hybrid screening system raises job offers by about 7% and reduces involuntary separations by 24% relative to the deterministic human-only screener and increases job offers by about 7% while reducing involuntary separations by 23% relative to the deterministic AI-only screener. In the high-turnover BPO environment we study, these changes are sizable: they imply that a system which already improves on human screening (AI-only) can itself be dominated by a hybrid design that exploits the complementarities between human and AI signals.

Since $\bar{I}$ is held fixed, these gains do not arise from a change in selectiveness, but instead reflect more precise screening: the firm both hires more workers and experiences fewer involuntary exits among those hires, consistent with a higher correlation between the posterior index B and true match quality T.

Taken together with our reduced-form evidence that AI screening alone improves offers and retention relative to human screening, these counterfactuals suggest a monotone pattern: moving from human to AI to a hybrid human–AI system yields progressively better matches, driven by the additional and more balanced information that the firm extracts from applicants.

4.2.2 Dimension-Specific Screeners

We next consider a system motivated directly by the preliminary estimates in Table 2: the human screener appears to be slightly more precise in the first dimension, while the AI screener is more precise in the second dimension. This suggests a dimension-specific design in which each screener is assigned exclusively to the dimension in which it performs best. Accordingly, the human screener provides a signal s_1^H about θ_1, the AI screener provides a signal s_2^{AI} about θ_2, and the firm also observes the language and analytical tests with precision κ. Given these inputs, the firm forms posterior beliefs about the composite index according to:

$$\mathbb{E}[I \mid s_1^H, s_2^{AI}, LS, AS] = \alpha \left(\frac{\kappa t_1 + \rho_1^H s_1^H}{1 + \kappa + \rho_1^H} \right) + (1 - \alpha) \left(\frac{\kappa t_2 + \rho_2^{AI} s_2^{AI}}{1 + \kappa + \rho_2^{AI}} \right),$$

and the firm again employs the worker whenever this posterior exceeds the threshold $\bar{I}$. Unconditional job offers and retention rates are computed as before.

Relative to each stand-alone screener, the dimension-specific system generates smaller but still economically meaningful gains. Job offers increase by about 1% and involuntary separations fall by about 2% relative to the deterministic human-only screener; compared to the deterministic AI-only screener, job offers rise by about 1% and involuntary separations fall by about 1%. These gains are modest compared to those in the hybrid system because the firm observes only one interview signal per dimension rather than two, even though each signal now comes from the screener with higher precision on that dimension. The modest gains are a consequence of the preliminary estimates: both human and AI precisions are already relatively high and differ only slightly, so assigning dimension specific screeners yields limited incremental improvement. We expect this exercise to be more informative in the full estimation. Indeed, as the setup changes, for example when we repeat the calculation under the specification in which 25% of voluntary departures are treated as low-quality, the corresponding improvements rise to about 5%. In short, depending on the final estimates, meaningful gains may arise from assigning each screener to the dimension it measures most accurately.

Another benefit of using dimension-specific screeners emerges once we account for the costs of conducting interviews with humans and AI. Even when two screeners have similar precision in a given dimension, differences in their employment costs can make specialization advantageous. For instance, suppose AI outperforms humans in screening the second dimension, while both perform similarly on the first. If employing AI—particularly in the early stages—is more expensive than employing humans, then our analysis suggests an approach in which AI is used exclusively for the second dimension, while humans handle the first. This strategy allows firms to benefit not only from improved overall precision but also from lower total costs. Thus, the gains considered here represent only part of the broader picture.

More broadly, these preliminary counterfactuals illustrate the potential complementarities between human and AI screeners. Rather than viewing AI as a one-for-one substitute for human judgment, the model suggests that integrating both technologies—either by combining their signals or by specializing them across dimensions—can improve match quality and retention. In the high-turnover BPO environment we study, improvements in retention at a common hiring threshold $\bar{I}$ translate directly into lower training and replacement costs, so even modest percentage changes are economically meaningful.

However, we emphasize that these quantitative effects are based on preliminary parameter estimates and should be interpreted as suggestive; the comprehensive estimation

described in Section 4.3 will provide final estimates and formal standard errors. Once we estimate the full model in Section 4.3, we will revisit these designs with fully estimated parameters and extend them to counterfactuals that explicitly involve choice, information structures, and interview sequencing.

4.3 Comprehensive Estimation Procedure

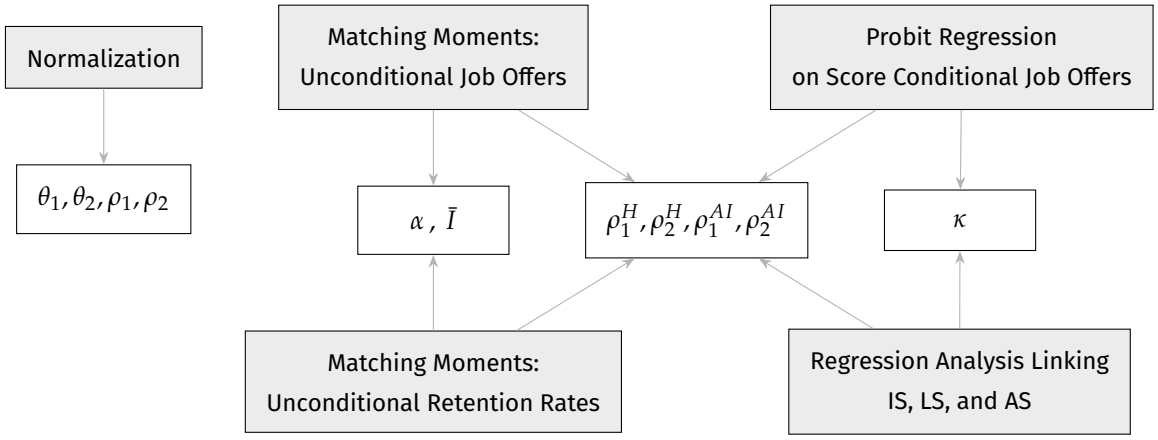
As outlined above, the approach combines estimation and calibration exercises, providing a proof of concept that the model is identifiable from the available data. The next step is to estimate all model parameters using a joint maximum-likelihood procedure that incorporates both hiring and retention outcomes, conditional on candidates' test and interview scores. This unified estimation strategy allows us to recover all parameters simultaneously and compute their associated standard errors. The overall identification and estimation strategy is summarized in Figure 7: Panel 7a outlines the identification of the primitives, and Panel 7b illustrates how welfare differences between the deterministic and choice-based screening systems are used to infer the information structure under which the industry operates.

Moreover, as the model section makes clear, the introduction of choice may have heterogeneous implications: it can benefit applicants, the firm, both—-or, in some cases, help one side while harming the other. As shown there, the nature of these effects depends critically on whether the firm incorporates the informational content of choice into its inference, and, if so, whether workers are aware that such inference takes place. The theoretical framework specifies how the model's parameters integrate welfare changes for both firms and workers under these alternative screening systems.

This stage of the estimation exploits the differences in recruitment and retention between the three interviewer systems in the randomized sample of eligible 67,056 applications described in Section 2.3, comparing the deterministic Human-AI systems with the choice-based system in which applicants endogenously select their interviewer. The subsequent stage of the estimation therefore leverages differences between the hiring and retention decisions observed under the choice treatment and those under the deterministic treatments (where applicants were directly assigned an AI or Human screener) to infer (i) whether the firm uses choice as an informative signal, (ii) whether workers recognize this behavior, and (iii) what the resulting equilibrium outcomes are.

Figure 7: Identification Estimation Strategy

(a) Identification of technological parameters.



(b) Information Structure Identification: Choice vs Deterministic Screening.

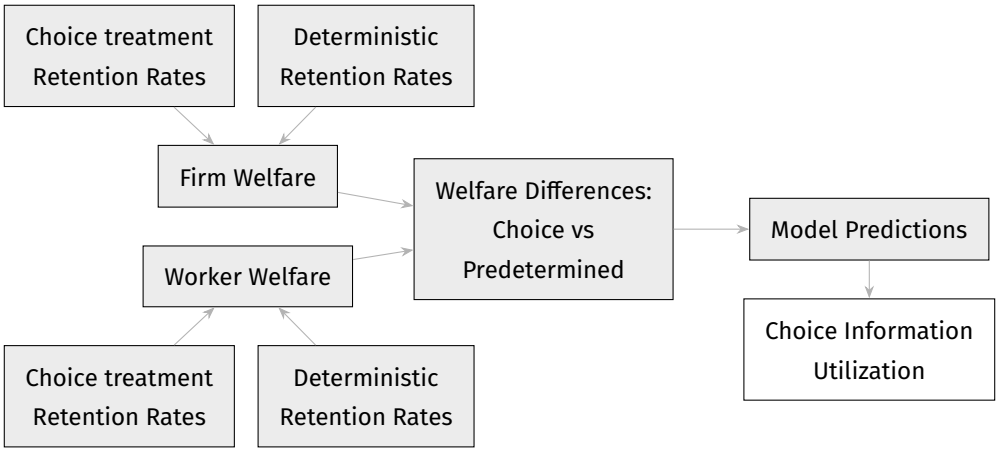


Figure 7 summarizes this overall identification and estimation strategy. First, Panel 7a illusrates the first step of the comprehensive estimation. Normalization pins down the distribution of both applicant-type dimensions, as described in Step 1 of Section 4.1. Probit regressions on interview, language, and analytical scores—together with regressions of interview scores on language and analytical scores—provide additional moments, as in Steps 2 and 3. Unconditional job offers and unconditional retention rates deliver two further sets of moments, as in Steps 4 and 5. Taken together, these components identify the key technological parameters governing screener precision, firm selectiveness, the relative importance of each dimension, and test accuracy.

Second, Panel 7b shows how we use the choice-based screening system to identify behavior and welfare. Conditional on the technological parameters, the model maps job offers and retention outcomes in the deterministic Human-AI systems and in the choice-

based screening system into firm and worker welfare. Comparing welfare across systems under alternative information structures—for example, whether the firm uses candidates' screener choices as an additional signal and whether workers anticipate this—yields distinct predictions for the observed welfare gap between Choice and predetermined screening. As emphasized in Section 3, the primitive parameters remain fixed across information structures, but the implied welfare differences vary with the information structure. Matching these welfare differences in the data then allows us to identify the information structure consistent with equilibrium behavior.

4.4 Additional Counterfactual Screening Systems

Based on these estimates, we can revisit counterfactual exercises in Section 4.2 with fully estimated parameters and standard errors, and extend them to screening systems that explicitly involve choice, information design, and interview sequencing. Below we name a few.

Choice Versus Deterministic Systems First, we will compare the welfare performance of the Choice system with that of the deterministic human-only and AI-only systems, using the estimated model to construct firm, worker, and social welfare measures under each assignment rule. This exercise will quantify (i) how much of the welfare gain from AI adoption comes from deploying AI per se versus granting applicants autonomy over the interviewer, and (ii) whether self-selection into AI versus human interviews moves hiring and retention closer to, or further away from, the welfare-maximizing assignment.

Choice Versus Specialized Human-AI Assignment Second, we will contrast the Choice system with a specialized design in which humans and AI are assigned to different dimensions, allowing us to evaluate whether specialization on language versus analytical skills can outperform offering applicants a choice of screener. Using the estimated precisions (ρ_j^H, ρ_j^{AI}) and the joint distribution of types, we can compute the welfare of an optimally specialized assignment rule and compare it to Choice, thereby measuring the welfare loss from decentralizing assignment to applicants rather than letting the firm exploit the relative comparative advantages of human and AI screeners across dimensions and identifying which subsets of workers (e.g., language-strong vs analytically-strong types) gain or lose under specialization.

We also plan to examine a design with two screener options constructed to maximize the informativeness of applicants' choices. For example, one option may assign humans to the language dimension and AI to the analytical dimension, while the other reverses

these roles, aiming to maximize the information conveyed by the choice.

Choice as a Screening Signal and Information Design Third, we will compare alternative information structures that govern if and how the firm can condition on applicants' screener choices. Specifically, we will contrast screening systems in which (i) the firm does not observe whether a choice was made, (ii) the firm observes screener choice and is allowed to treat it as an additional signal of type, and (iii) applicants are explicitly informed or not informed that their choice enters the screening rule. Within the estimated model, this will allow us to quantify the incremental predictive and welfare value of using choice as a signal, the welfare loss from ignoring it, and how transparency about this use of choice affects selection and the division of surplus between firms and workers.

Double Interviewing and Order Effects Fourth, we will analyze a double-interview design in which each candidate is seen by both a human and an AI screener, in randomized order. Within our framework, this corresponds to allowing effective precisions (ρ_j^H, ρ_j^{AI}) and/or the hiring threshold $\bar{I}$ to depend on whether an AI or a human interview occurs first, in the spirit of the order effects documented by Radbruch and Schiprowski (2025). This exercise will quantify the incremental value of combining human and AI signals in a single decision rule, measure how order effects distort inferred types and hiring decisions, and compare the welfare of double-interviewing to that of single-interview and Choice systems.

These analyses will enable a richer understanding of the equilibrium and welfare implications of introducing choice into the hiring process and will contribute to identifying the optimal way to integrate AI into hiring decisions. In doing so, the results will inform both theory and practice regarding the design of efficient screening mechanisms, clarifying when and how human and algorithmic judgment should be combined to achieve better outcomes for firms and applicants alike.

5 Conclusion

We study a setting in which applicants choose between two screening technologies—human or AI—and analyze three environments reflecting different hiring practices: one where firms ignore the information conveyed by choice, one where firms internalize it but applicants do not anticipate this, and one where both sides recognize that in equilibrium screener choice is informative and will be used in hiring. These environments yield distinct welfare implications. Allowing choice can harm the firm—when it fails to internalize

the information conveyed by choice—strictly benefit the firm—when it internalizes the information but applicants do not anticipate this—or improve firm outcomes when the screeners specialize in different skill dimensions, as in the equilibrium case. For workers, the information conveyed by choice generally helps high-ability applicants. It typically harms low-ability workers, who benefit only if the firm ignores the informational content of choice. Thus, whether offering choice benefits one party, the other, or both depends critically on how the information embedded in choice is used.

We complement the model section with a preliminary estimation and calibration exercise that identifies key parameters and enables counterfactual analyses. We quantify the gains from having *both* human and AI interviews (rather than only one). We then assess the value of *specialized screening*, where each screener focuses on the dimensions it measures best. We also laid the groundwork for a full structural estimate, which will allow us to quantify the practical value of offering choice, determine which theoretical environment best matches industry behavior, compare choice to adding more interview rounds, and evaluate whether firms update beliefs appropriately.

Our results shed light on how human and AI screeners complement one another and on the role of choice in labor-market screening. We show that screener choice—initially conceived as an applicant protection measure—can benefit both firms and applicants. Rather than posing a trade-off between AI and humans, our findings point to the more nuanced question of where AI should be deployed and where human judgment remains essential. This perspective, we hope, can guide the design of hybrid systems that combine technological advances with human expertise.

References

Acemoglu, Daron, and Pascual Restrepo (2018). "Artificial intelligence, automation, and work." In *The economics of artificial intelligence: An agenda*. University of Chicago Press, 197–236. [5]

Acemoglu, Daron, and Pascual Restrepo (2019). "Automation and new tasks: How technology displaces and reinstates labor." *Journal of Economic Perspectives* 33 (2): 3–30. [5]

Acemoglu, Daron, and Pascual Restrepo (2022). "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90 (5): 1973–2016. [5]

Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz (2023). "Combining human expertise with artificial intelligence: Experimental evidence from radiology." Working paper. National Bureau of Economic Research. [5]

Agarwal, Nikhil, Alex Moehring, and Alexander Wolitzky (2025). "Designing Human-AI Collaboration: A Sufficient-Statistic Approach." Working Paper, First version: April 18, 2025. AEA Registry AEARCTR-0013990. [6]

Aka, Ada, Emil Palikot, Ali Ansari, and Nima Yazdani (2025). "Better Together: Quantifying the Benefits of AI-Assisted Recruitment." Working Paper. [6]

Ambuehl, Sandro, B. Douglas Bernheim, and Axel Ockenfels (2021). "What Motivates Paternalism? An Experimental Study." *American Economic Review* 111 (3): 787–830. [5]

Angelova, Victoria, Will S. Dobbie, and Crystal Yang (2023). "Algorithmic Recommendations and Human Discretion." NBER Working Paper 31747. National Bureau of Economic Research. [6]

Ash, Elliott, Soumitra Shukla, and Jason Sockin (2025). "Interviews." Unpublished manuscript. [6]

Athey, Susan C., Kevin A. Bryan, and Joshua S. Gans (2020). "The Allocation of Decision Authority to Human and Artificial Intelligence." *AEA Papers and Proceedings* 110: 80–84. [5]

Autor, David, and Neil Thompson (2025). "Expertise." *NBER Working Paper*, (w33941): [5]

Autor, David H., and David Scarborough (2008). "Does Job Testing Harm Minority Workers? Evidence from Retail Establishments*." *Quarterly Journal of Economics* 123 (1): 219–77. [10]

Benzell, Seth G, Laurence J Kotlikoff, Guillermo LaGarda, and Jeffrey D Sachs (2015). "Robots are us: Some economics of human replacement." Working paper. National Bureau of Economic Research. [5]

Berg, Jeff, Avinash Das, Vinay Gupta, and Paul Kline (2018). "Smarter Call-Center Coaching for the Digital World." Technical Report. New York: McKinsey & Company. [8]

Bernheim, B. Douglas (2016). "The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics." *Journal of Benefit-Cost Analysis* 7 (1): 12–68. [5]

Bernheim, B. Douglas, and Antonio Rangel (2007). "Toward Choice-Theoretic Foundations for Behavioral Welfare Economics." *American Economic Review* 97 (2): 464–70. [5]

Bernheim, B. Douglas, and Antonio Rangel (2009). "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics*." *Quarterly Journal of Economics* 124 (1): 51–104. [5]

Bonney, Kathryn, Cory Breaux, Catherine Buffington, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, and Keith Savage (2024). "Tracking Firm Use of AI in Real Time: A

Snapshot from the Business Trends and Outlook Survey." NBER Working Paper 32319. National Bureau of Economic Research. [5]

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond (2025). "Generative AI at Work." *Quarterly Journal of Economics* 140 (2): 889–942. [5, 6]

Brynjolfsson, Erik, and Andrew McAfee (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & company. [5]

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What can machines learn and what does it mean for occupations and the economy?" In *AEA papers and proceedings*. vol. 108, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 43–47. [5]

Buesing, Eric, Vinay Gupta, Sarah Higgins, and Raelyn Jacobson (2020). "Customer Care: The Future Talent Factory." Technical Report. New York: McKinsey & Company. [8]

Bushong, Benjamin, Carolina Corral, Xavier Giné, Aprajit Mahajan, and Enrique Seira (2025). "Autonomy and Technology Adoption." Working Paper. [5]

Chakraborty, Ishita, Khai Chiong, Howard Dover, and K. Sudhir (2024). "Can AI and AI-Hybrids Detect Persuasion Skills? Salesforce Hiring with Conversational Video Interviews." *Marketing Science* 44 (1): 30–53. [6]

Cowgill, Bo (2019). "Bias and Productivity in Humans and Machines." Research Paper. Columbia Business School. Available at SSRN: https://ssrn.com/abstract=3584916. [6]

Cowgill, Bo, Pablo Hernández-Lagos, and Nataliya Wright (2024). "Does AI Cheapen Talk? Theory and Evidence from Global Entrepreneurship and Hiring." Working Paper. [6]

Czarnitzki, Dirk, Gastón P. Fernández, and Christian Rammer (2023). "Artificial Intelligence and Firm-Level Productivity." *Journal of Economic Behavior & Organization* 211: 188–205. [5]

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani (2023). "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality." *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013): [5]

Estrada, Ricardo (2019). "Rules versus Discretion in Public Service: Teacher Hiring in Mexico." *Journal of Labor Economics* 37 (2): 545–79. [6]

Frankel, Alex (2021). "Selecting Applicants." *Econometrica* 89 (2): 615–45. [6]

Galdin, Anaïs, and Jesse Silbert (2025). "Making Talk Cheap: Generative AI and Labor Market Signaling." Job Market Paper. [6]

Hernandez, Edrin (2024). "Analyzing BPO Employment Statistics Philippines: Trends and Insights." Technical Report. Magellan Solutions. [8]

Hoffman, Mitchell, Lisa Kahn, and Danielle Li (2018a). "Discretion in Hiring." *Quarterly Journal of Economics* 133 (2): 765–800. [6]

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li (1, 2018b). "Discretion in Hiring." *The Quarterly Journal of Economics* 133 (2): 765–800. [10]

Hoffman, Mitchell, and Christopher Stanton (2024). "People, Practices, and Productivity: A Review of New Advances in Personnel Economics." Working paper w32849. National Bureau of Economic Research. [6]

Hoong, Ruru, and Bnaya Dreyfuss (2025). "Calibrated Coarsening: Designing Information for AI-Assisted Decisions." *Available at SSRN 5286198*, [6]

Iansiti, Marco, and Karim R. Lakhani (2020). *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Boston, MA: Harvard Business Review Press. [5]

Jabarian, Brian, and Luca Henkel (2025). "Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews." Working paper. University of Chicago Booth School of Business, and Erasmus University Rotterdam. Available at SSRN: https://ssrn.com/abstract=5395709. [1–3, 6, 7, 10, 11, 66, 83]

Jovanovic, Boyan (1979). "Job Matching and the Theory of Turnover." *Journal of Political Economy* 87 (5, Part 1): 972–90. [10]

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan (2018). "Human Decisions and Machine Predictions." *Quarterly Journal of Economics* 133 (1): 237–93. [5]

Li, Danielle, Lindsey Raymond, and Peter Bergman (2025). "Hiring as Exploration." *Review of Economic Studies*, (6): [6]

McDaniel, Michael A., Deborah L. Whetzel, Frank L. Schmidt, and Steven D. Maurer (1994). "The Validity of Employment Interviews: A Comprehensive Review and Meta-Analysis." *Journal of Applied Psychology* 79 (4): 599–616. [6]

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia S. Foster, and Nikolas Zolas (2024). "AI Adoption in America: Who, What, and Where." *Journal of Economics & Management Strategy* 33 (2): 375–415. [5]

Otis, Nicholas, Rowan Clarke, Solene Delecourt, David Holtz, and Rembrand Koning (2024). "The uneven impact of generative AI on entrepreneurial performance." [5]

Parasurama, Prasanna, and Panos Ipeirotis (2025). "When Good Screeners Go Bad: A Correlation-Driven Paradox in Two-Stage Selection Pipelines and Human-AI Collaboration." Working Paper. [6]

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer (2023). "The impact of ai on developer productivity: Evidence from github copilot." *arXiv preprint arXiv:2302.06590*, [5]

Radbruch, Jonas, and Amelie Schiprowski (2025). "Interview sequences and the formation of subjective assessments." *Review of Economic Studies* 92 (2): 1226–56. [6, 45, 51]

Sallaz, Jeffrey J. (2019). *Lives on the Line: How the Philippines Became the World's Call Center Capital*. New York, NY: Oxford University Press. [8]

Stevenson, Megan T, and Jennifer L Doleac (2024). "Algorithmic risk assessment in the hands of humans." *American Economic Journal: Economic Policy* 16 (4): 382–414. [6]

Vravosinos, Orestis (2025). "Multidimensional screening of strategic candidates." Working Paper. [6]

Wiles, Emma, and John J Horton (2024). "More, but worse: The impact of ai writing assistance on the supply and quality of job posts." *Mass. Inst. Technol.(MIT) Sloan*, [6]

Yang, Frank, Piotr Dworczak, and Mohammad Akbarpour (2023). "Comparison of Screening Devices." Working Paper. [6]

Appendix

A Additional Analysis

A.1 Supporting Calculations

Supporting calculation for Observation 1(i).

Step 1. Monotonicity of the basic fractions. Assume $\rho_1^{AI} > \rho_1^H$ and $\rho_2^{AI} > \rho_2^H$. By symmetry, the same argument applies if both inequalities are reversed. For $a > 0$, set $f(x) = \frac{x}{(a+x)^2}$, with derivative $f'(x) = \frac{a-x}{(a+x)^3}$. Hence f decreases for $x > a$. Since $\rho_i^m > \rho_i$, we obtain

$$\frac{\rho_1^{AI}}{(\rho_1+\rho_1^{AI})^2} < \frac{\rho_1^H}{(\rho_1+\rho_1^H)^2}, \qquad \frac{\rho_2^{AI}}{(\rho_2+\rho_2^{AI})^2} < \frac{\rho_2^H}{(\rho_2+\rho_2^H)^2},$$

and thus $R_{AI} < R_H$. Similarly, $g_i(x) = \frac{x}{\rho_i+x}$ is strictly increasing, so

$$\frac{\rho_1^{AI}}{\rho_1+\rho_1^{AI}} > \frac{\rho_1^H}{\rho_1+\rho_1^H}, \qquad \frac{\rho_2^{AI}}{\rho_2+\rho_2^{AI}} > \frac{\rho_2^H}{\rho_2+\rho_2^H}.$$

Therefore, $R_{AI} < R_H$.

Step 2. Comparing A_{AI} with A_H. We have

$$A_{AI} = \frac{\alpha}{R_{AI}} \cdot \frac{\rho_1^{AI}}{\rho_1+\rho_1^{AI}}, \qquad A_H = \frac{\alpha}{R_H} \cdot \frac{\rho_1^H}{\rho_1+\rho_1^H}.$$

Thus

$$\frac{A_{AI}}{A_H} = \frac{\frac{\rho_1^{AI}}{\rho_1+\rho_1^{AI}}}{\frac{\rho_1^H}{\rho_1+\rho_1^H}} \cdot \frac{R_H}{R_{AI}}.$$

The first factor is > 1 by monotonicity of g_1, and the second factor is > 1 since $R_{AI} < R_H$. Therefore $\frac{A_{AI}}{A_H} > 1$, i.e.

$$A_{AI} > A_H. \tag{5}$$

Step 3. Comparing B_{AI} with B_H. Analogously,

$$B_{AI} = \frac{1-\alpha}{R_{AI}} \cdot \frac{\rho_2^{AI}}{\rho_2+\rho_2^{AI}}, \qquad B_H = \frac{1-\alpha}{R_H} \cdot \frac{\rho_2^H}{\rho_2+\rho_2^H},$$

so

$$\frac{B_{AI}}{B_H} = \frac{\frac{\rho_2^{AI}}{\rho_2+\rho_2^{AI}}}{\frac{\rho_2^{H}}{\rho_2+\rho_2^{H}}} \cdot \frac{R_H}{R_{AI}}.$$

Both factors are > 1, so we have

$$B_{AI} > B_H. \tag{6}$$

Step 4. Sign of B. By (5), the denominator $A_{AI} - A_H > 0$. By (6), the numerator $B_H - B_{AI} < 0$. Therefore

$$B = \frac{B_H - B_{AI}}{A_{AI} - A_H} < 0.$$

□

Supporting calculations for Observation 1(ii). **Step 1. Definition and monotonicity of the basic fractions.** Note that

$$\frac{A_{AI}}{A_H} = \frac{\frac{\rho_1^{AI}}{\rho_1+\rho_1^{AI}}}{\frac{\rho_1^{H}}{\rho_1+\rho_1^{H}}} \cdot \frac{R_H}{R_{AI}}, \qquad \text{and} \qquad \frac{B_H}{B_{AI}} = \frac{\frac{\rho_1^{H}}{\rho_1+\rho_1^{H}}}{\frac{\rho_1^{AI}}{\rho_1+\rho_1^{AI}}} \cdot \frac{R_{AI}}{R_H}.$$

Assume $\rho_1^{AI} > \rho_1^H$ and $\rho_2^{AI} < \rho_2^H$. A symmetric argument applies if the inequalities are reversed. We either have

$$\frac{\rho_1^{AI}}{(\rho_1+\rho_1^{AI})^2} < \frac{\rho_2^{AI}}{(\rho_2+\rho_2^{AI})^2} \quad \text{or the opposite.}$$

$$\frac{\partial}{\partial\alpha}\left(\frac{R_{AI}}{R_H}\right) = \frac{2(1-\alpha)\alpha\left(\frac{\rho_1^{AI}\rho_2^H}{(\rho_1+\rho_1^{AI})^2(\rho_2+\rho_2^H)^2} - \frac{\rho_1^H\rho_2^{AI}}{(\rho_1+\rho_1^H)^2(\rho_2+\rho_2^{AI})^2}\right)}{\left(\frac{\alpha^2\rho_1^H}{(\rho_1+\rho_1^H)^2} + \frac{(1-\alpha)^2\rho_2^H}{(\rho_2+\rho_2^H)^2}\right)^2} < 0$$

The inequality follows because, for any α, the sign of the derivative is determined by

$$\frac{\rho_1^{AI}\rho_2^H}{(\rho_1+\rho_1^{AI})^2(\rho_2+\rho_2^H)^2} - \frac{\rho_1^H\rho_2^{AI}}{(\rho_1+\rho_1^H)^2(\rho_2+\rho_2^{AI})^2}.$$

Note that

$$\frac{d}{dx}\left(\frac{xy}{(\rho_1+x)^2(\rho_2+y)^2}\right) = \frac{(\rho_1-x)y}{(\rho_1+x)^3(\rho_2+y)^2},$$

$$\frac{d}{dy}\left(\frac{xy}{(\rho_1+x)^2(\rho_2+y)^2}\right)=\frac{(\rho_2-y)x}{(\rho_2+y)^3(\rho_1+x)^2}.$$

Replacing x with ρ_1^{AI} or ρ_1^H, and y with ρ_2^H or ρ_2^{AI}, the derivative is negative since $\rho_1^{AI},\rho_1^H>\rho_1$ and $\rho_2^{AI},\rho_2^H>\rho_2$. Moreover, because $\rho_1^{AI}>\rho_1^H$ and $\rho_2^H>\rho_2^{AI}$, the difference is always negative.

Consequently, $\frac{A_{AI}}{A_H}$ ($\frac{B_H}{B_{AI}}$) is monotonically increasing (decreasing) in α.

Step 2. Comparing extreme points of the functions. Next, note that

$$\lim_{\alpha\to 0}\frac{A_{AI}}{A_H}=\underbrace{\frac{\rho_1^{AI}\left(\rho_1+\rho_1^H\right)\left(\rho_2+\rho_2^{AI}\right)\sqrt{\frac{\rho_2^H}{\rho_2^{AI}}}}{\left(\rho_1+\rho_1^{AI}\right)\rho_1^H\left(\rho_2+\rho_2^H\right)}}_{A(0)},\quad \lim_{\alpha\to 1}\frac{A_{AI}}{A_H}=\underbrace{\sqrt{\frac{\rho_1^{AI}}{\rho_1^H}}}_{A(1)},$$

$$\lim_{\alpha\to 0}\frac{B_H}{B_{AI}}=\underbrace{\sqrt{\frac{\rho_2^H}{\rho_2^{AI}}}}_{B(0)},\quad \lim_{\alpha\to 1}\frac{B_H}{B_{AI}}=\underbrace{\frac{\sqrt{\frac{\rho_1^{AI}}{\rho_1^H}}\left(\rho_1+\rho_1^H\right)\left(\rho_2+\rho_2^{AI}\right)\rho_2^H}{\left(\rho_1+\rho_1^{AI}\right)\rho_2^{AI}\left(\rho_2+\rho_2^H\right)}}_{B(1)}.$$

Thus, $\frac{A_{AI}}{A_H}$ increases monotonically from $A(0)$ to $A(1)$, while $\frac{B_H}{B_{AI}}$ decreases monotonically from $B(0)$ to $B(1)$. For suitable parameter values, $A(0)$ and $B(1)$ can be below 1, so β can be either positive or negative.

A sufficient condition for $A(0)>1$ is

$$\frac{\rho_1^{AI}}{\rho_1^H}\cdot\frac{\rho_1+\rho_1^H}{\rho_1+\rho_1^{AI}}>\sqrt{\frac{\rho_2^H}{\rho_2^{AI}}}.$$

A sufficient condition for $B(1)>1$ is

$$\sqrt{\frac{\rho_1^{AI}}{\rho_1^H}}\cdot\frac{\rho_1+\rho_1^H}{\rho_1+\rho_1^{AI}}>1.$$

□

Supporting Calculations for Observation 2(i).
By definition, $\hat{\theta}_j=\mathbb{E}[\theta_j\mid\mathcal{I}]$. Hence, using by the law of iterated expectations,

$$\mathbb{E}\left[\hat{\theta}_j-\theta_j\right]=\mathbb{E}\left[\mathbb{E}[\theta_j\mid\mathcal{I}]-\theta_j\right]=\mathbb{E}\left[\mathbb{E}[\theta_j\mid\mathcal{I}]-\mathbb{E}[\theta_j\mid\mathcal{I}]\right]=0.$$

The conditional version follows immediately:

$$\mathbb{E}\left[\hat{\theta}_j - \theta_j \,\middle|\, \mathcal{I}\right] = \mathbb{E}[\theta_j \mid \mathcal{I}] - \theta_j$$

has conditional expectation

$$\mathbb{E}\left[\mathbb{E}[\theta_j \mid \mathcal{I}] - \theta_j \,\middle|\, \mathcal{I}\right] = \mathbb{E}[\theta_j \mid \mathcal{I}] - \mathbb{E}[\theta_j \mid \mathcal{I}] = 0.$$

Thus the posterior mean is unbiased, both unconditionally and (in the sense of mean-zero error) conditional on the information set that includes the selection event. □

Lemma 1. *Let $Z \sim \mathcal{N}(0,1)$ and define the (upper–tail) inverse Mills ratio*

$$\lambda_+(\gamma) \;=\; \frac{\varphi(\gamma)}{1-\Phi(\gamma)} \;=\; \mathbb{E}[\,Z \mid Z > \gamma\,].$$

Then, for every $\gamma \in \mathbb{R}$,

$$\lambda_+(\gamma) \;>\; \gamma, \qquad \textit{hence} \qquad f(\gamma) := \lambda_+(\gamma)^2 - \gamma\,\lambda_+(\gamma) = \lambda_+(\gamma)\big(\lambda_+(\gamma) - \gamma\big) \;>\; 0.$$

Proof. Since $Z \mid Z > \gamma$ is supported on (γ, ∞) and has a strictly positive density on every interval $(\gamma, \gamma + \varepsilon)$, we have

$$\mathbb{E}[\,Z \mid Z > \gamma\,] \;\geq\; \gamma \cdot 1 + \varepsilon \cdot \mathbb{P}(Z > \gamma + \varepsilon \mid Z > \gamma) \;>\; \gamma \quad \text{for any } \varepsilon > 0,$$

so $\lambda_+(\gamma) > \gamma$. Because also $\lambda_+(\gamma) > 0$, the product $\lambda_+(\gamma)\big(\lambda_+(\gamma) - \gamma\big)$ is strictly positive, proving $f(\gamma) > 0$. □

Supporting Calculations for Observation 2(ii).

Assume without loss that $m^* = AI$ and $A_{AI} > A_H$; symmetric arguments follow for all other cases. Then

$$\mathrm{Var}\big(\theta_1 \mid s_1^m, s_2^m, m = AI\big) = \frac{1}{\rho_1 + \rho_1^m} \;-\; \frac{1}{(\rho_1 + \rho_1^m)^2}\, f(\gamma)\, \sigma_L^2,$$

$$\mathrm{Var}\big(\theta_2 \mid s_1^m, s_2^m, m = AI\big) = \frac{1}{\rho_2 + \rho_2^m} \;-\; \frac{\beta_1^2}{(\rho_2 + \rho_2^m)^2}\, f(\gamma)\, \sigma_L^2,$$

with

$$f(\gamma) = \left(\frac{\varphi(\gamma)}{1-\Phi(\gamma)}\right)^2 - \gamma\, \frac{\varphi(\gamma)}{1-\Phi(\gamma)}.$$

By Lemma 1, $f(\gamma) > 0$ for all $\gamma \in \mathbb{R}$. Therefore

$$\mathrm{Var}\big(\theta_j \mid s_1^m, s_2^m, m = AI\big) \;\leq\; \frac{1}{\rho_j + \rho_j^m} \qquad (j = 1, 2),$$

□

A.2 The Non-Selection Benchmark Further Analysis

A.2.1 Single Dimension Case $(J = 1)$

With a single dimension

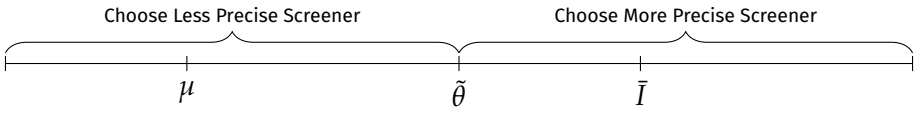
$$p_m(\theta) = \Pr\big(I^m \geq \bar{I} \mid \theta\big) = \Pr\big(\hat{\mu}^m \geq \bar{I} \big| \theta\big) = 1 - \Phi\left(\frac{\rho + \rho^m}{\sqrt{\rho^m}}\left(\bar{I} - \frac{\rho\,\mu + \rho^m\,\theta_i}{\rho + \rho^m}\right)\right).$$

Since $\Phi(\cdot)$ is a monotonically increasing function of the argument, the applicant will choose the screener with the lower argument. Assume without loss that $\rho^{AI} > \rho^H$. We set the arguments equal and solve for the θ value that would make the applicant indifferent between the two screeners.

$$\frac{\rho + \rho^A I}{\sqrt{\rho^A I}}\left(\bar{I} - \frac{\rho\mu + \rho^A I\,\tilde{\theta}}{\rho + \rho^A I}\right) = \frac{\rho + \rho^H}{\sqrt{\rho^H}}\left(\bar{I} - \frac{\rho\mu + \rho^H\,\tilde{\theta}}{\rho + \rho^H}\right) \quad \Longrightarrow\; \tilde{\theta} = \bar{I} - \frac{(\bar{I} - \mu)\rho}{\sqrt{\rho^{AI}\rho^H}}$$

$$m^*(\theta) = \begin{cases} \text{AI}, & \text{if } \theta \geq \tilde{\theta} \\ H, & \textit{otherwise} \end{cases}$$

Figure A.1: Single Dimensional Choices



Notes: There exists a cutoff type $\tilde{\theta}$ such that applicants with types below $\tilde{\theta}$ prefer the noisier screener, as the added noise gives them a chance to surpass the threshold $\bar{I}$. In contrast, applicants with types above $\tilde{\theta}$ prefer the more precise screener.

Some Comparative Statics. The cutoff $\tilde{\theta}$ responds to changes in the model's parameters in intuitive ways. First, $\frac{\partial\tilde{\theta}}{\partial\bar{I}} > 0$: as the firm's hiring bar $\bar{I}$ rises, the threshold above which applicants prefer the more precise screener also increases. By contrast, $\frac{\partial\tilde{\theta}}{\partial\rho} < 0$, so that

greater prior precision lowers the cutoff. The cutoff is increasing in the prior mean, $\frac{\partial \tilde{\theta}}{\partial \mu} > 0$, and its behavior with respect to signal precision ρ^m follows similar logic. Finally, in the limiting case where the high-precision screener becomes fully revealing, $\lim_{\rho^{AI} \to \infty} \tilde{\theta} = \bar{I}$, meaning that applicants would choose it only if their type is at least $\bar{I}$.

A.2.2 Single Dimension Case $(J = 1)$ (Continuous Adjustment)

With a single dimension

$$p_m(\theta) = \Pr\left(I_i^m \geq \bar{I} \mid \theta\right) = \Pr\left(\hat{\mu}_i^m \geq \bar{I} \middle| \theta_i\right) = 1 - \Phi\left(\frac{\rho + \rho^m}{\sqrt{\rho^m}}\left(\bar{I} - \frac{\rho\,\mu + \rho^m\,\theta_i}{\rho + \rho^m}\right)\right).$$

Since $\Phi(\cdot)$ is a monotonically increasing function of the argument, we focus on how the argument changes as the signal precision ρ^m increases.

$$\frac{\partial \frac{\rho+\rho^m}{\sqrt{\rho^m}}\left(\bar{I} - \frac{\rho\,\mu+\rho^m\,\theta_i}{\rho+\rho^m}\right)}{\partial \rho^m} = \frac{(\rho^m - \rho)}{2(\rho^m)^{3/2}}\left(\bar{I} - \frac{\rho\mu + \rho^m\theta_i}{\rho + \rho^m}\right) + \frac{\rho + \rho^m}{\sqrt{\rho^m}}\frac{\rho(\mu - \theta_i)}{(\rho + \rho^m)^2}$$

The first term, $\frac{(\rho^m - \rho)}{2x^{3/2}}$, determines how the precision of the estimator changes. If $\rho^m < \rho$, as the precision of the screener increases, more weight is loaded on it, this leads to a lower overall precision. However, once $\rho^m = \rho$, any further increase in the precision of the screener increases the overall precision. Whether this is good or bad, depends on whether the expected value is above or below $\bar{I}$, as captured by $\left(\bar{I} - \frac{\rho\mu+\rho^m\theta_i}{\rho+\rho^m}\right)$. Setting the above equal to 0 and solving for θ_i yields

$$\tilde{\theta} = \frac{\bar{I}(\rho^m - \rho) + \mu\rho}{\rho^m}$$

$$\lim_{\rho^m \to \infty} \tilde{\theta} \to \bar{I} \qquad \lim_{\rho \to 0} \tilde{\theta} \to \bar{I}$$

In both cases, in the limit, the realized signal fully determines the posterior, thus, the applicant would prefer more precision only if their type exceeds $\bar{I}$.

B Simulation Procedure

The following steps outline the simulation procedure in detail.

1. **Specify model parameters.**
    - Set prior means (μ_1, μ_2) and precisions (ρ_1, ρ_2).

- Specify screener signal precisions (ρ_1^m, ρ_2^m) for $m \in \{H, AI\}$.
- Choose the firm's weighting parameter $\alpha \in (0,1)$ and the hiring cutoff $\bar{I}$.
- Fix N, the number of applicant types θ to simulate.
- Initiate adjustments $\gamma_H = \gamma_{AI} = 0$.

2. **Sample applicant types and screener signals.**
   - For each applicant $i = 1, \ldots, N$ and dimension $j = 1, 2$, draw types independently as $\theta_{i,j} \sim \mathcal{N}(\mu_j, 1/\rho_j)$.
   - For each screener $m \in \{H, AI\}$, applicant $i = 1, \ldots, N$, and dimension $j = 1, 2$, draw signals independently as $s_{i,j}^m \sim \mathcal{N}(\theta_{i,j}, 1/\rho_j^m)$.
3. **Outer Loop: Compute the separating boundary, and screener choice.**
   - For each screener m, compute $R_m,\ A_m,\ B_m$, and afterward $\tilde{\beta}_0(\gamma_H, \gamma_{AI}),\ \beta_1$.
   - Given $\tilde{\beta}_0(\gamma_H, \gamma_{AI}),\ \beta_1$, for each participant, based on their θ values, choose screener $m = H$ or $m = AI$.
   - **Inner Loop: Grid Search** (τ_H, τ_{AI})
     - For each applicant $i = 1, \ldots, N$ based on their signals $s_{i,j}^m$ from their chosen screener m, compute the adjusted posterior composite index

$$\alpha \left(\frac{\rho_1 \mu + \rho_1^m s_{i,1}^m}{\rho_1 + \rho_1^m} \right) + (1 - \alpha) \left(\frac{\rho_2 \mu + \rho_2^m s_{i,2}^m}{\rho_2 + \rho_2^m} \right) + \tau_m$$

     - Compute hiring decisions based on if the above clears $\bar{I}$ or not.
     - Compute $W^F(\tau_H, \tau_{AI})$.
     - Update τ_H and τ_{AI} to the ones that maximized $W^F(\tau_H, \tau_{AI})$.
   - **Iterate until convergence** Repeat until $|\tau_j^{new} - \tau_j^{old}| < \varepsilon$. For some ε tolerance.
4. **Output and diagnostics.**
   - Record the final equilibrium boundary and equilibrium adjustments.
   - Calculate and plot firm, worker, and other conditional welfare outcomes.

C Estimation Related Analysis

C.1 Relevant Regressions

Table C.1 presents the regression of the human interview score on the language and analytical scores, as outlined in Step 2 of Section 4.1.

Table C.1: Predicting Interview Scores using Test Scores: Human Case

Dependent variable:	Standardized Interview Score
Language test score (std.)	0.184***
	(0.017)
Analytical test score (std.)	0.049***
	(0.018)
Observations	2,522

Note: *p<0.1; **p<0.05; ***p<0.01

Table C.2 presents the regression of the AI interview score on the language and analytical scores, again, as outlined in Step 2 of Section 4.1.

Table C.2: Predicting Interview Scores using Test Scores: AI Case

Dependent variable:	Standardized Interview Score
Language test score (std.)	0.147***
	(0.009)
Analytical test score (std.)	0.041***
	(0.009)
Observations	10,435

Note: *p<0.1; **p<0.05; ***p<0.01

Table C.3 presents the probit regression of the employment decision under the deterministic Human treatment, on interview and test scores, as outlined in Step 2 of Section 4.1. Here, we report the two parameters of interest for our estimation procedure, namely the coefficients on the language and analytical scores.

Table C.3: Predicting Job Offer Decisions using Test Scores: Human Case

Dependent variable:	Job Offer Made
Language test score (std.)	0.317***
	(0.030)
Analytical test score (std.)	0.089***
	(0.029)
Observations	2,522

Note: *p<0.1; **p<0.05; ***p<0.01

Similarly, Table C.4 presents the probit regression of the employment decision under the deterministic AI treatment, as outlined in Step 2 of Section 4.1.

Table C.4: Predicting Job Offer Decisions using Test Scores: AI Case

Dependent variable:	Job Offer Made
Language test score (std.)	0.405***
	(0.016)
Analytical test score (std.)	0.092***
	(0.015)
Observations	10,435

Note: *p<0.1; **p<0.05; ***p<0.01

C.2 Alternative Preliminary Estimated Parameters

Table C.5 presents the counterpart of Table 2 under the specification $v = 0.75$. In this case, we assume that among the voluntary departures, 75% would have met the firm's desired standards, while the remaining 25% would have fallen below those standards despite leaving voluntarily.

Table C.5: Preliminary Estimated Parameters: $v = 0.75$

$\hat{\bar{I}}$	$\hat{\alpha}$	$\hat{\rho}_1^H$	$\hat{\rho}_2^H$	$\hat{\rho}_1^{AI}$	$\hat{\rho}_2^{AI}$	$\hat{\kappa}$
0.89	0.29	2.66	2.52	2.57	3.16	0.32

Notes: The estimated parameters above result from following the steps outlined in Section 4.1, which combine estimation and calibration exercises. These parameters are subject to revision.

Similarly, Table C.5 presents the counterpart of Table 2 under the specification $v = 0.50$. In this case, we assume that among the voluntary departures, 50% would have met the firm's desired standards, while the remaining 50% would have fallen below those standards despite leaving voluntarily.

Table C.6: Preliminary Estimated Parameters: $v = 0.50$

$\hat{\bar{I}}$	$\hat{\alpha}$	$\hat{\rho}_1^H$	$\hat{\rho}_2^H$	$\hat{\rho}_1^{AI}$	$\hat{\rho}_2^{AI}$	$\hat{\kappa}$
0.78	0.25	0.79	0.75	0.9	1.1	0.32

Notes: The estimated parameters above result from following the steps outlined in Section 4.1, which combine estimation and calibration exercises. These parameters are subject to revision.

D AI voice agent

System overview. The AI interviewer used in our setting is a cloud-based voice agent, "Anna AI", that conducts live job interviews over the phone. It integrates three standard modules of modern conversational AI: multilingual automatic speech recognition (ASR), a job-specific large language model (LLM) that selects questions and follow-ups, and streaming text-to-speech (TTS) that converts generated text into synthetic speech. These components are orchestrated in real time so that the agent can ask questions, listen to the applicant, and respond with follow-ups in a natural conversational flow. The underlying interview script mirrors the human guideline described in the main text, so the human and AI interviewers implement the same sequence of topics and information requests.

Technical architecture. More concretely, the partner's AI voice agent is supplied by a specialized AI–SaaS vendor and stacks three generative technologies in real time: (i) multilingual ASR, (ii) a job-specific LLM, and (iii) streaming TTS. The ASR front end converts accented Filipino–English audio into text; the transcript is routed through an LLM fine-tuned on thousands of proprietary interviews, which selects the next question or follow-up; the response is rendered in natural speech by a neural vocoder and sent back to the candidate. End-to-end latency averages roughly 0.4–1.8 seconds, which makes turn-taking effectively indistinguishable from a human conversation because the entire stack runs on the vendor's dedicated GPU cluster. "Contextual pathway prompts" inject the relevant job description at every turn, while guardrails veto off-topic or non-compliant content.

Operational challenges. Deploying such an AI interviewer in live hiring interviews magnifies several classic challenges of language-model-based systems. The agent must remain coherent and on-topic throughout the conversation, avoiding hallucinations or irrelevant digressions; it must also be robust to attempts by applicants to "game" the interview by reading from scripts or repeating buzzwords. Spoken interactions add further constraints: the agent must cope with background noise, variation in accents and speaking pace, and tight latency requirements, since delays above about a second break conversational flow and can raise stress or dissatisfaction on the applicant side. Meeting these constraints required custom acoustic models, controller-layer guardrails, and low-latency GPU inference, deployment investments beyond an off-the-shelf API call. In this paper, we take the resulting AI interviewer as a given screening technology and focus on how its use should be designed relative to human interviewers.

E Detailed recruitment process

E.1 Engagement-Score Pre-Screening

In the remote recruitment mode, applicants pass through an engagement-based pre-screening before being queued for interview scheduling. Each applicant receives a base score that depends on how their profile entered the system (self-application versus recruiter-sourced). The algorithm then adjusts this score based on simple interaction signals: valid phone and e-mail information, successful delivery of SMS and e-mail messages, openings of those messages, and explicit "Yes" or "No" replies to outreach. Positive signals raise the score, and negative replies reduce it; missing contact details do not change the score. Once the cumulative score crosses a low, account-specific threshold, the applicant is moved into the interview queue; applicants who never reach the threshold are not contacted. PSG deliberately sets these thresholds low in order to interview as many candidates as possible, so this pre-screen is lightly selective.[18] Randomization into the Human, AI, and Choice screener regimes occurs after this step, at the applicant level within the firm's recruitment system.

[18]The firm's exact point values, thresholds, and queuing code are proprietary; see Jabarian and Henkel (2025) for further discussion.

E.2 Invitation text

E.2.1 Treatments *Human interviewer* and *AI interviewer*

Email: If a recruiter decides to interview an applicant, the following invitation text is sent via email to them. The text is identical across the *Human interviewer* and *AI interviewer* conditions.

[Subject line:] Interview Invitation: Schedule Your interview for the [Name Position]

Hi [First name] [Last name],

I hope you're doing well! My name is [Name recruiter], from [Name recruiting partner], the recruiting partner of [Name recruiting firm]. We've had a chance to review your application for the [Name Position]. We currently have an immediate need to fill this position. Not all roles require [Position-specific] experience, providing opportunities for various backgrounds and skill sets. If you are interested, please click on the button below.

[Button "Get Interview Call"]

[Clicking the button will redirect applicants to the firm's website, where they can schedule their interview.]

Phone: At the same time as the email is sent, the following invitation text is sent via text message to applicants' phones. Again, the text is identical across the *Human interviewer* and *AI interviewer* conditions.

Subject: Interview Invitation: Schedule Your Interview for the [Name position]

Hi [First name] [Last name],

We are reaching out to you regarding your application for the [Name position]. We'd like you to schedule your interview. We've sent the interview invitation to your email.

E.2.2 Treatment *Choice interviewer*

Applicants in the *Choice interviewer* treatment receive the same email and phone message as in the other two treatments. The only difference is that clicking the button "Get Interview Call" in the email will redirect them to an interview scheduling preference page. The text on the page is as follows:

Interview Scheduling Preference

Congratulations! You have been shortlisted for an interview. Please select your preferred interviewer:

- *AI Interviewer: The call can be scheduled at your convenience.*
- *Human Interviewer: You'll need to schedule the interview based on the human recruiter's availability.*

[After selecting the interviewer, applicants can schedule their interview exactly as in the other two treatments.]

F Additional figures

Figure F.1: Recruitment process

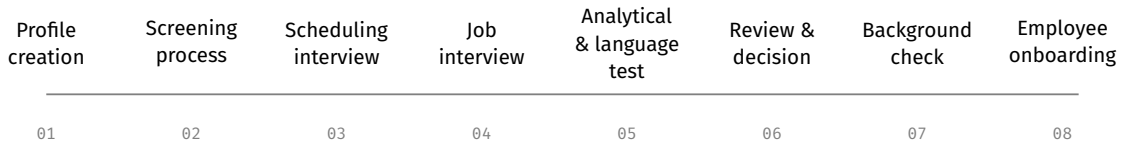
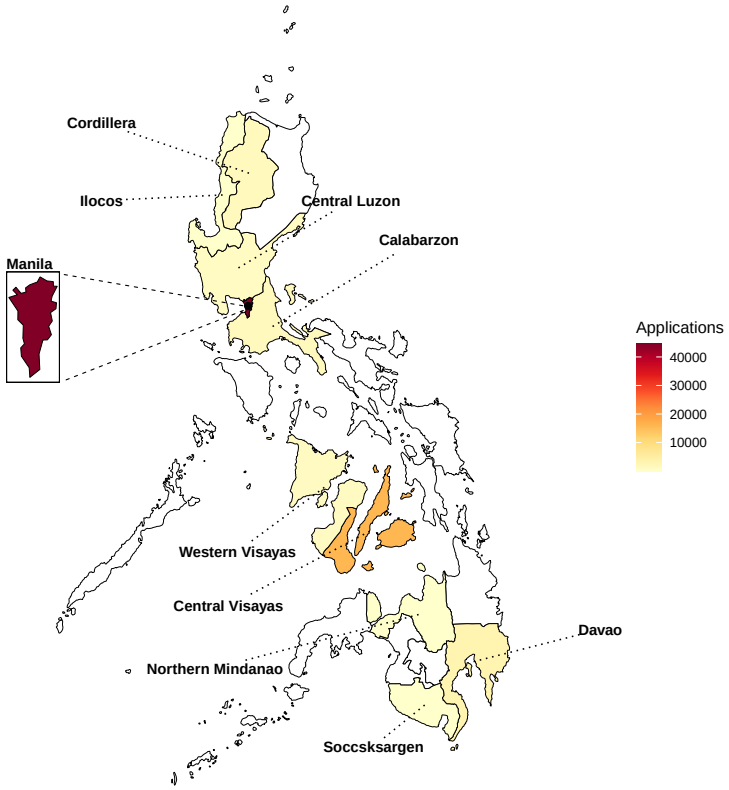


Figure F.2: Distribution of applications by recruitment center location



G Job description wording

Title: Customer Expert

Our *Customer Service Representatives* and *Technical Support Representatives* are vital members of our company. You will field customer inquiries and find innovative ways to respond. You will have the chance to work in a highly collaborative and engaging environment that provides dynamic work experience with different cultures, as well as unlimited opportunities to grow your potential and develop your career.

As a *Customer Service Representative / Technical Support Representative*, your responsibilities will include:

- Handling and carefully responding to all customer inquiries via inbound calls and email
- Providing excellent customer service through active listening
- Working with confidential customer information in a secure manner
- Aiming to resolve issues on the first call by being proactive
- Appropriately and adequately communicating with customers

Working hours:
Monday to Friday – 8:00 PM to 5:00 AM PH Time / 10:00 PM to 7:00 AM PH Time

Background Requirements:

- NBI clearance
- Birth certificate
- Fit-to-work clearance

Compensation: Depends on job, between Php 16,000 and Php 25,000

Minimal Requirements:

- SHS Grad / HS Grad
- Average communication skills

H Examples of full interviews

The following transcripts present two sample interviews, one from the *Human Interviewer* and one from the *AI Interviewer* condition. As noted in the main text, the content of interviews varies across application sites and job profiles. These examples are therefore provided for illustrative purposes and are not necessarily representative of all interviews conducted in the experiment. To protect applicant confidentiality, the transcripts have been lightly edited to preserve anonymity and prevent any possibility of direct or indirect identification.

H.1 Human Interviewer example

Interviewer: Hello.
Candidate: Hello, good morning.
Interviewer: Hi, this is [Anonymized] from [Anonymized] Recruitment. I'll be conducting the interview and I'll be asking you personal details and some common interview questions. Then I'll be giving you feedback prior to your position, okay? [Topic 1: Introduction]
Candidate: Yes, of course.
Interviewer: Now, I just want to remind you that this phone call is recorded for quality assurance purposes.
Candidate: Yes.
Interviewer: Okay. So can you state your complete name, your middle name, your age, and your birthplace?
Candidate: My name is [Anonymized] and I'm [Anonymized].
Interviewer: I'm sorry, how old are you?
Candidate: [Anonymized].
Interviewer: [Anonymized]?
Candidate: Yeah, correct.
Interviewer: Can you turn down on your phone maybe?
Candidate: Yes.
Interviewer: Have you ever worked for [Anonymized] in the past?
Candidate: No. This is my second job in [Anonymized].
Interviewer: And where are you currently residing? [Topic 2: Location Verification]
Candidate: In [Anonymized].
Interviewer: If you are ever hired, will there be any issues commuting or traveling from your place that you work?

Candidate: No, there's no issues at all because right now we do have a direct train from [Anonymized] going to your place.
Interviewer: May I know what will be your mode of transportation or how long will it take you to travel? [Topic 2: Mode of Transportation]
Candidate: Yeah, just a few. I think it will be around an hour.
Interviewer: Can you state your highest educational attainment? [Topic 3: Educational Background]
Candidate: I did [Anonymized].
Interviewer: Hello, I'm sorry. I'm sorry I can't hear you clearly. Can you turn the muffled sound if you can? Can you transfer for a better reception or maybe using your headset?
Candidate: Yeah, I'm currently using my headset right now.
Interviewer: If you can, can we just mute your phone so that you can hear me clearly?
Candidate: Okay, ill just remove my headset. Sorry. Hello?
Interviewer: Hello, can you hear me clearly?
Candidate: Okay, that's better. Thank you so much. First, I graduated [Anonymized]. Then I took up my masters degree in [Anonymized] but unfortunately I just finished [Anonymized] because there were some circumstances that I have to face here in the Philippines. Thats why I go back there.
Interviewer: Okay. May I know what's your masters degree?
Candidate: Masters degree in [Anonymized], majoring in [Anonymized].
Interviewer: May I ask you if you have plans to pursue your studies? [Topic 4: Candidate's Commitment and Attrition Risks - Future Study Plan]
Candidate: Right now, no because I need to focus first with [Anonymized]. Thats the reason why I came back here in the Philippines.
Interviewer: Thank you. Are you currently employed?
Candidate: Sorry?
Interviewer: Are you currently employed?
Candidate: No. I'm not employed right now.
Interviewer: Do you have any pending applications from the other companies? [Topic 4: Candidate's Commitment and Attrition Risks - Other Applications]
Candidate: To be honest, I just submitted my application and fortunately this is the first company who called me for an interview. Thats why I'm here right now.
Interviewer: Thank you. How about any previous work experience? All in all? Any foreign? [Topic 5: Past Employment History]
Candidate: First, my first three years I worked in [Anonymized] as [Anonymized]. Then after that, I went abroad to [Anonymized]. I worked as a [Anonymized] for [Anonymized]

years. Then I go back here in the Philippines to work as [Anonymized] for [Anonymized]. Then I went to [Anonymized] and worked as an [Anonymized].
Interviewer: So the most recent job that you had is in [Anonymized], right?
Candidate: Yeah, correct.
Interviewer: Can you please tell me more about it?
Candidate: As a [Anonymized], I'm the one responsible for [Anonymized]. All the problems related with [Anonymized], they put in [Anonymized]. Then every morning I have to check [Anonymized] and check all the priorities which need to handle first. Also, I have to give feedback to [Anonymized]. And I also receive some [Anonymized] from the customers regarding the problems or issues with their equipment or any computer related problems. And especially, I need to resolve all those issues to reduce the downtime on their side. Basically, that's all my job.
Interviewer: And may I know what's the reason why you left that job?
Candidate: I left my job. Thats what I'm saying because I have to face some problems here in the Philippines because [Anonymized], they have some issues. I have to weigh in which is better for the future. But I will not sacrifice my [Anonymized]. So I'd rather choose [Anonymized] rather than [Anonymized].
Interviewer: May I ask if you have any certificates of employment from the previous job?
Candidate: Yeah, I do.
Interviewer: So what made you decide to apply?
Candidate: Yeah, because I submitted applications and fortunately this is the first company who called me for an interview. Also, I think since I'm challenging, like trying something like [Anonymized] which is very in demand in this country. And I also want to improve my comm skills as well because working in [Anonymized] will improve your comm skills.
Interviewer: May I know what's your salary expectation? [Topic 6: Compensation]
Candidate: To be honest, I would like to. I dont have a specific since I dont have any background experience in [Anonymized]. So I would like to assume that from the range between [Anonymized] to [Anonymized]
Interviewer: [Anonymized] to [Anonymized], right?
Candidate: Yes, yes.
Interviewer: The range that we can offer you if ever you're going to get hired will be from [Anonymized] to [Anonymized] depending on what account are you going to be with. Are you okay with it?
Candidate: Yes.
Interviewer: Okay. So I'm sure you have [Anonymized] experience.

Candidate: Yeah, [Anonymized]. The customers will call [Anonymized] to help to resolve their problem or some issues. For example, for instance, they have an internet issue and the client or the customer will call [Anonymized] and the job of [Anonymized] is to resolve, to find the root cause and to resolve the issues of [Anonymized].
Interviewer: And how would you define quality [Anonymized]?
Candidate: Oh yeah. First, you must be polite to the service and you must also handle the average customer and you should be able to resolve all the [Anonymized] issues and concerns. Especially with less supervision from [Anonymized]. You must be able to handle that on your own.
Interviewer: So if ever youre going to encounter irate or difficult [Anonymized], how do you think you can handle them effectively?
Candidate: Yeah, I always stay calm. I already encountered with irate [Anonymized] in [Anonymized]. I stay calm and explain to them everything in [Anonymized]: I explain everything clearly and I try to focus on the problem and solution to resolve the issues of [Anonymized].
Interviewer: What are skills do you think you have that can be a great asset for the job position?
Candidate: First, I'm a hardworking and dynamic self-disciplined person and I can also work as an individual or a team player and I'm good at handling [Anonymized].
Interviewer: If ever you're going to encounter conflict between your peers or your colleagues, how are you going to address them?
Candidate: I will have to help them personally and privately, not in front of the audience, and help them and try to resolve the issues.
Interviewer: Do you have any [Anonymized] or at least most of the time I mentioned?
Candidate: No. [Anonymized]
Interviewer: Do you have any other certifications aside from your work?
Candidate: No, not at all.
Interviewer: Do you have any travel plans within the next six months, international or local? [Topic 7: Candidate's Commitment and Attrition Risks - Planned Vacation]
Candidate: No, I don't have one.
Interviewer: Have you ever had the willingness to work on-site?
Candidate: Yeah, of course.
Interviewer: You can work in shifting schedules, in extended hours, holidays, weekends, weekdays? [Topic 7: Candidate's Commitment and Attrition Risks - Working during Holidays]
Candidate: Yeah, no problem at all. Because when I worked in [Anonymized] for [Anonymized],

I worked in various shifts for at least [Anonymized].
Candidate: So working in various shifts is not a problem at all.
Interviewer: Can you start as soon as possible if ever youre going to get hired?
Candidate: Yeah, absolutely. And I'm willing to start.
Interviewer: And may I know how long are you planning to work here with us? [Topic 7: Candidate's Commitment and Attrition Risks - Long-term Intention]
Candidate: Sorry?
Interviewer: How long are you planning to work here with us?
Candidate: As long as I'm enjoying my job in our company, Ill pursue my goal. Because if I'm going to start here as a regular [Anonymized], thats not my goal.
Candidate: My goal is to become a higher position, which I can use my masters degree in handling people, something like that.
Interviewer: Many of you already have your government documents like [Anonymized] to help you by any chance? [Topic 8: Data Verification]
Candidate: Right now I dont have [Anonymized] yet, but I do have [Anonymized].
Interviewer: And may I know what valid IDs do you have with you?
Candidate: I do have my [Anonymized].
Interviewer: And are you vaccinated?
Candidate: [Anonymized]
Interviewer: Do you have any medical issues that we need to know?
Candidate: [Anonymized].
Interviewer: Okay, I think that concludes the end of our interview. The next step is the assessment. Please wait for the contact officer to call you, okay? [Topic 9: Conclusion]
Candidate: Okay, thank you so much.

H.2 AI Interviewer example

Interviewer: Hi [Anonymized]. This is Anna, [Anonymized] AI recruiter and I am calling about the [Anonymized] role you applied to recently. Do you have fifteen minutes to talk about your application? [Topic 1: Introduction + Candidate Availability]
Candidate: Yes.
Interviewer: Great! Since I am an AI recruiter, as I ask you questions, if you are not clear on my question, please feel free to ask me for clarification. Does that work for you?
Candidate: Yes.
Interviewer: Ok, I also want to let you know a human recruiter will review the recording from our discussion today and will make the final decision on your application for

employment. The questions I will be asking you are the same questions my human counterpart would ask. Does that sound OK?
Candidate: Yes.
Interviewer: Let's get started. Have you ever worked for [Anonymized] in the past?
Candidate: No. This is my first time working in, um, [Anonymized].
Interviewer: That's okay, we get a lot of new applicants. Can you tell me a little bit about yourself and why you're interested in this role?
Candidate: Firstly, good afternoon. I'm [Anonymized] from [Anonymized] and I am an [Anonymized] graduate and have also completed [Anonymized] semester of [Anonymized] at [Anonymized] in [Anonymized]. My year of experience in [Anonymized] has equipped me with strong customer service and skills, and I'm confident in my ability to make a significant contribution to your team.
Interviewer: That's great to hear about your experience and education. This position is an on-site position and is located at [Anonymized]. Can you share with me where you reside? [Topic 2: Location Verification]
Candidate: I'm residing in [Anonymized]. Is there any available position in [Anonymized]?
Interviewer: Our office is actually located in [Anonymized]. How far a drive would this be for you?
Candidate: I think, um, it's one and a half hour from [Anonymized]. But it would be better if there's any open position in [Anonymized] .
Interviewer: Ok, how about transportation to and from work. Can you share with me the mode of transportation you would use to commute to our office daily? [Topic 2: Mode of Transportation]
Candidate: Oh, from [Anonymized], I'd estimate my commute to [Anonymized] would be around [Anonymized] minutes to [Anonymized] depending on traffic. But to ensure punctuality, uh, also plan to commute at least [Anonymized] before my shift, starts since it's a one ride trip only, and I'm confident about reaching the time reaching the office on time.
Interviewer: The minimum age requirement for this position is 18 years of age, are you at least 18 years of age? [Topic 3: Age Requirement]
Candidate: I'm [Anonymized] years old.
Interviewer: Most of the opportunities we have are on a shifting schedule where you could work during holidays and/or weekends. Have you had previous employment that has required you to work a similar schedule? [Topic 4: Candidate's Commitment and Attrition Risks - Working during Holidays]
Candidate: Yes.

Interviewer: In addition to a shifting schedule, we sometimes require overtime. Would you be able to accommodate extended hours when needed? [Topic 4: Candidate's Commitment and Attrition Risks - Working Overtime]
Candidate: Yes. I will.
Interviewer: I'd like to ask you some questions about your career goals and motivations. When you read our job advertisement, what attracted you to the position at [Anonymized] [Topic 4: Candidate's Commitment and Attrition Risks - Motivations]
Candidate: Actually, my previous coworker referred me to because as per him, there's open roles for a [Anonymized] account and as I have [Anonymized] experience and I think I will be fit for this role. That's why I applied for this job.
Interviewer: That's interesting. As you consider your career goals, how long could you envision working for [Anonymized] and performing the duties of this job? [Topic 4: Candidate's Commitment and Attrition Risks - Long-term Intentions]
Candidate: I think I see myself um, still working to your company [Anonymized] to [Anonymized] years. If there's a good position or good environment, I think I I'm able to work as long as I work on my previous job.
Interviewer: As you consider making a job change, what is important to you as you evaluate job opportunities?
Candidate: Right now, my goal is to work at your company since I as I check on my multimedia reviews on [Anonymized] that your company's team good. That that's why I want to focus on working with company as well.
Interviewer: Thank you for your answers, I appreciate the insight.
Candidate: That's why I think I didn't need to find another job. I'm just interested in the job position that you have open.
Interviewer: I'd like to shift back to a few more questions about the position we are considering you for. Our training requires mandatory attendance. If selected for the position, would you be able to attend all training?
Candidate: Yes. I will. Mm-mm.
Interviewer: Do you have any planned vacations over the next 6 months? [Topic 4: Candidate's Commitment and Attrition Risks - Planned Vacations]
Candidate: No. I don't have any plans. My, um, go for this year or this month is work. Focus on work.
Interviewer: Are there any activities that may affect your attendance or any prior commitments that you have to attend to?
Candidate: For now, there aren't any.
Interviewer: Let's take a few minutes to review your education background. What is your

highest level of education? [Topic 5: Educational Background]
Candidate: I am a [Anonymized] graduate and have also completed two semester of [Anonymized] at [Anonymized].
Interviewer: That's interesting. Are there any activities that may affect your attendance or any prior commitments that you have to attend to?
Candidate: For now, I think there are no activities in the future that affect my attendance because I will now talk as to work, and I have no any vacation pending.
Interviewer: Let's take a few minutes to review your education background. You've already told me about your highest level of education completed, which is [Anonymized] and having completed [Anonymized] at [Anonymized]. I just want to confirm, is that correct?
Candidate: Yes.
Candidate: I'm an [Anonymized] and a [Anonymized], um, completer, uh, for [Anonymized]. I think it's correct.
Interviewer: Since you have completed [Anonymized], do you plan to return to [Anonymized]?
Candidate: For now, I don't think so. I just my goal for now is to get a job. Permanent job, to sustain my bill needs.
Interviewer: If we can shift our discussion to compensation expectations, that would be helpful. What is your expected salary? [Topic 6: Compensation]
Candidate: I consider my ex considering my experience in customer facing enroll, I was hoping for a salary closer to [Anonymized] to [Anonymized].
Interviewer: The roles that we have open right now are between the salary package range of [Anonymized] to [Anonymized] per Month, depending on your qualifications, you could be offered a salary package between these amounts. Would you be okay with this range?
Candidate: Is there any closer for [Anonymized] to [Anonymized]? May I know the specific or may I know the account first?
Interviewer: The roles that we have open right now are between the salary package range of [Anonymized] to [Anonymized], depending on your qualifications, you could be offered a salary package between these amounts. Would you be okay with this range?
Candidate: Yes. I would be okay.
Interviewer: Thank you again for the information you've shared. I have a few more questions to go over, so let's continue. I'd like to talk about your [Anonymized] experience. can you share with me the volume of [Anonymized] on a daily basis? [Topic 7: Past Employment History]
Candidate: My previous job, I have been in [Anonymized] near [Anonymized] and I re-

ceived [Anonymized] well, we received [Anonymized] about [Anonymized]. That's our goal for the day, and I get [Anonymized] that's passing the [Anonymized] for a day.
Interviewer: That sounds like a significant vo-
Candidate: yes. Continue.
Interviewer: Were these [Anonymized] solving technical issues or were you responding to [Anonymized]?
Candidate: We conduct troubleshooting and also we create [Anonymized], and we talk to a customer and explain their needs, what their needs, and also like that.
Interviewer: Now, I'd like to ask about your previous work experience. Have you worked with other companies before this role?
Candidate: My previous job, I ended my journey there [Anonymized].I think. Yeah. That's my last job.
Interviewer: So, it sounds like you have some work experience, specifically from your last job.
Candidate: Yes.
Interviewer: Now, could you please provide the name of the most recent company you worked for?
Candidate: I've been for [Anonymized] for a year, and my first job was [Anonymized]. Then the second one, is I've I have [Anonymized].
Interviewer: Now, can you tell me what are your employment dates for those companies you mentioned.
Candidate: My employment date like, probably this job I've been there for [Anonymized] and ended [Anonymized]. It's almost [Anonymized]
Interviewer: So, your employment dates are from [Anonymized] to [Anonymized]. That's helpful to know.
Candidate: Mhmm. That's correct.
Interviewer: Now. Can you tell methe location of the company we are discussing. Where is this company located?
Candidate: This company must my previous job was located at [Anonymized].
Interviewer: So your previous job was located at [Anonymized]. Now, I'd like to ask, do you have any subsequent work experience with other companies?
Candidate: I'm also I also work as health care account at [Anonymized]. Before I apply for at [Anonymized], then I applied for this.
Interviewer: That's interesting. Now, could you please provide the name of the second most recent company you worked for?
Candidate: Second reason, company I work for. I work for [Anonymized]. I'm [Anonymized]

there. I've been there almost [Anonymized].
Interviewer: I'd like to know more about your work history. What are your employment dates for these previous roles, if you have them available?
Candidate: For a while, let me double check here because I don't I don't remember. I think it's to be sure, let's me let me double check here on my resume for a while.
Interviewer: Where is this company located?
Candidate: Okay. Before I applying for a [Anonymized] as a [Anonymized] for [Anonymized] account. Before that, I've been in a [Anonymized] as a [Anonymized]. Duration was [Anonymized]. It's located at [Anonymized].
Interviewer: So, it seems like you have some experience in customer service. Do you have any subsequent work experience with other companies besides [Anonymized] and [Anonymized]?
Candidate: I also work at [Anonymized]. I've been there for duration is [Anonymized] to [Anonymized] as a [Anonymized]. All my job all my previous job experience is [Anonymized] roles.
Interviewer: Could you please provide the name of the third most recent company you worked for?
Candidate: My third is the [Anonymized]. It's a [Anonymized], and my position there a [Anonymized]. The duration is [Anonymized] until [Anonymized]. This company was located at [Anonymized].
Interviewer: You mentioned you worked at [Anonymized] as a [Anonymized], can you tell me why you left that company?
Candidate: I left at [Anonymized] because I'm sick.
Interviewer: If you are offered the job, when would you be able to start? [Topic 8: Availability]
Candidate: I can for now because since I don't have any work, I can work
Interviewer: Do you have any pending applications this this week in [Anonymized]
Candidate: This is my first application before I ended my contract on my previous job.
Interviewer: Do you have any pending applications with other companies?
Candidate: I don't have any pending applications.
Interviewer: What is your contact number? [Topic 9: Data Verification]
Candidate: Okay. My contact number and my cell phone number [Anonymized]
Interviewer: Could y
ou please share your [Anonymized] number with me?
Candidate: Okay. Same as I said earlier, my phone number through [Anonymized] is [Anonymized] **Interviewer:** May I know your Marital/Civil Status, please?

Candidate: [Anonymized].
Interviewer: [Anonymized] do you have any questions you would like for me to document and have reviewed by the recruiting team?
Candidate: Yes. I want to ask if there's any available position for your local account in tour site at [Anonymized] because I was prepared my by my previous by my previous coworker and I really want to join their team because I I believe, like, I have a role because or I fitted the job this job because I have a [Anonymized]. As per my referral, your current employee, that's health care account for your local account or bill care account is now open for hiring.
Interviewer: I have noted your question, and I will make sure to document it, after this call, someone from our team will get back to you regarding the available position for [Anonymized], and we will also verify the information about [Anonymized] being open for hiring.
Candidate: Okay. Thank you. That's all.
Interviewer: Thank you [Anonymized]. This concludes the questions I have for you today. Following this interview, our recruiting team will be reviewing the summary of our conversation to determine next steps regarding your application for employment. Due to the high demand and urgency to fill our openings, we encourage you to complete an online assessment. By taking this assessment right away, you'll demonstrate your commitment to this opportunity. [Topic 10: Conclusion + Assessment Test]
Candidate: Mhmm.
Interviewer: Great! We'll see you on the next hiring process. Good luck!

I Additional tables

Table I.1: Client firm industries

Client firm sector	Percentage (%) $(n = 41)$
Technology	22.00
Finance	14.60
Transportation	14.60
Tourism	9.80
Healthcare	7.30
Insurance	7.30
Telecommunication	7.30
Hospitality	4.90
Manufacturing	4.90
Retail	4.90
Energy	2.40

Notes: Distribution of client firm industries for which the company is hiring customer service representatives.

Table I.2: Treatment balance tests

Variable	*Human Interviewer* (1)	*AI Interviewer* (2)	*Choice of Interviewer* (3)	H_0: (1) = (2) p-value (4)	H_0: (1) = (3) p-value (5)	H_0: (2) = (3) p-value (6)
Panel A: Full sample						
Gender (Women=1)	0.60	0.60	0.61	0.21	0.50	0.04
Source is referral	0.19	0.19	0.19	0.96	0.38	0.26
Source is digital ad	0.59	0.59	0.59	0.47	0.85	0.34
Mode is Walk-in	0.26	0.26	0.26	0.45	0.66	0.85
Initial engagement score	37.06	37.12	37.21	0.80	0.62	0.73
Observations	13,557	40,103	13,396			
Panel B: *Walk-in mode*						
Gender (Women=1)	0.56	0.56	0.56	1.00	0.84	0.86
Source is referral	0.12	0.11	0.11	0.54	0.62	0.36
Source is digital ad	0.05	0.06	0.06	0.74	0.70	0.55
Observations	3,478	10,421	3,469			
Panel C: *Remote mode*						
Gender (Women=1)	0.62	0.61	0.62	0.15	0.48	0.02
Source is referral	0.22	0.22	0.22	0.66	0.63	0.30
Source is digital ad	0.77	0.77	0.78	0.66	0.66	0.32
Initial engagement score	50.48	50.42	50.48	0.75	0.99	0.73
Observations	10,079	29,682	9,927			

Notes: Columns (1) - (3) display mean values of variables for the three treatments. "Source is referral" and "Source is digital ad" are binary variables equal to one if the applicant applied through a referral or a digital job advertisement, respectively. Columns (4) - (6) display p-values obtained using pairwise t-tests (variable "Initial engagement score") or proportion tests (all other variables).

J Research transparency

This paper uses de-identified administrative and test data from the natural field experiment described in Jabarian and Henkel (2025). We implement no new experimental interventions for the present study: all recruitment, randomization, and data collection were carried out as part of that project. The experiment and associated data collection were approved by the University of Chicago Institutional Review Board (protocols #IRB24-1894 and #IRB25-1002) and pre-registered on the AEA RCT Registry (trial #15385), where the pre-analysis plan is publicly available.[19] Our analysis focuses on a subset of their pre-registered variables and adds a structural model and counterfactual welfare analyses that were not specified in the pre-analysis plan.

[19] See https://www.socialscienceregistry.org/trials/15385.