

Discriminated but not Discouraged: Investment Responses to Algorithmic Discrimination*

Yung-Shiang Jasmine Yang[†]

Abstract

We study how individuals respond to algorithmic evaluation with group-based priors and whether disadvantage discourages or instead induces compensatory investment. In a pre-registered online experiment ($N = 553$), participants are randomly assigned to either an advantaged or disadvantaged identity group and choose how much to invest to improve their hiring prospects under an algorithm that applies different prior beliefs across groups. Participants in the disadvantaged group invest 17.7% more than advantaged participants, consistent with compensatory investment rather than disengagement. We then test two interventions designed to encourage investment: information-based “Role Model” treatments and a “No Risk” treatment that conditions investment repayment on hiring. While the role model treatments have limited effects, the “No Risk” treatment increases investment among disadvantaged participants by 24%. Our findings suggest that reducing downside exposure associated with investing may be more effective than social-comparison interventions in encouraging self-investment under algorithmic statistical discrimination.

JEL classification: C91, D91, J71

Keywords: algorithmic bias, statistical discrimination, unequal opportunity, investment behavior, algorithmic hiring, online experiment

*I thank Emmanuelle Auriol, Thomas Buser, Haochi Chen, Vessela Daskalova, Paolo Falco, Marco Fongoni, Diogo Geraldes, Emma N. Guirado Forment, Oliver Hauser, Aditi Kharb, Kevin Kloiber, Manuel E. Lago, Lucie Martin, Ivan Pastine, Deirdre Robertson, Mikhail Safronov, Margaret Samahita, Nikita Sangwan, Nora Strecker, Vasilisa Werner, Kun Wu, Bin Zhang, participants at the 2025 European Doctoral Group in Economics (EDGE) Conference, the University College Dublin School of Economics Research & Reading Group Seminar, the 2026 Irish Postgraduate and Early Career Economics (IPECE) Workshop, the 50th International Association for Research in Economic Psychology (IAREP) Conference, and the 3rd Annual Conference of the Network of Behavioural and Experimental Economists based in the UK (BEE UK) for detailed feedback, helpful suggestions and stimulating conversations. All study procedures of the experiment described were approved by the ethics committee at University College Dublin (081-HS-LR-25-Yang-Samahita). The experiment and hypotheses were pre-registered on AsPredicted under ID #243245 [<https://aspredicted.org/f8jk-5xns.pdf>]. Funding from University College Dublin School of Economics is gratefully acknowledged.

[†]School of Economics, University College Dublin, Belfield, Dublin 4, Ireland; E-mail: yung-shiang.yang@ucdconnect.ie

1 Introduction

Algorithms now play an increasingly prominent role in high-stakes decisions, especially in hiring.¹ All of the top 20 Fortune 500 companies rely on data-driven recruitment tools to assist in screening and evaluating candidates (Ajunwa & Greene 2019, Bigu & Cernea 2019). Employers such as Cisco, Target, and Ikea use algorithms trained on past applications to predict hiring success (Bogen & Rieke 2018, Zhang & Yencha 2022). While these technologies promise increased efficiency, consistency, and reduced taste-based discrimination (Houser 2019, Van Esch et al. 2019), they also raise growing concerns over fairness and unequal opportunity. In particular, researchers highlight how hiring algorithms may replicate or amplify historical inequalities by relying on group-level priors, a mechanism closely related to statistical discrimination (Arrow et al. 1973, Phelps 1972, Barocas & Selbst 2016, Binns 2018). Because algorithmic systems can be deployed at scale and applied consistently across large populations, such disparities may become persistent and widespread (Kordzadeh & Ghasemaghaei 2022, Awad et al. 2023), as illustrated by Amazon’s now-defunct hiring algorithm that penalized resumes containing the word “women” (Hamilton 2018, Vincent 2018).

An additional reason to study algorithmic evaluation is that individuals often perceive algorithmic decision-makers differently from human evaluators. Prior research documents skepticism and aversion toward algorithmic evaluators, concerns about fairness (Dineen et al. 2004, Araujo et al. 2020), and differences in how people morally evaluate discrimination by algorithms relative to humans (Bigman et al. 2023, Dargnies et al. 2024). These findings suggest that algorithmic evaluation constitutes a distinct environment in which individuals make decisions. Yet while a growing literature examines attitudes toward algorithmic evaluators, much less is known about how individuals adjust costly behavior when facing unequal algorithmic evaluation.

These differences matter because perceptions of unequal opportunity shape beliefs about economic mobility and attitudes toward inequality (McCall et al. 2017, Fehr et al. 2024). Theory and empirical evidence suggest that discriminatory evaluation can generate self-fulfilling prophecies, where disadvantaged individuals rationally under-invest, thereby reinforcing initially unfavorable beliefs (Lundberg & Startz 1983, Coate & Loury 1993, Fryer et al. 2005, Fang & Moro 2011, Glover et al. 2017, Patty & Penn 2023). Yet other evidence documents the opposite response: individuals sometimes increase effort in response to perceived disadvantage, a phenomenon known as the “over-efforting” or “twice-as-hard” effect (Caputo 2002, 2007, Evans et al. 2019, Isik et al. 2021, Gutman & Younas 2025). While both mechanisms have been studied extensively in human-mediated environments, much less is known about how individuals respond when unequal evaluation is implemented by algorithmic decision-makers.

To answer these questions, we conduct a pre-registered incentivized online experiment ($N = 553$) that simulates a stylized hiring process with algorithmic screening. Participants are randomly assigned to either a disadvantaged Blue group or an advantaged Green group and, independently, to a latent productivity type (“top performer” or not) with equal probability (50/50), which is unobservable to both the participant and the algorithm. Participants then decide how much to invest in improving their prospects. Higher investment increases their probability of becoming a “top performer,” thereby raising their likelihood of achieving higher observable pre-employment test scores.

The algorithm does not observe true productivity directly. Instead, it forms hiring decisions based on pre-employment test scores combined with group identity, using group-specific priors

¹We use hiring algorithm as the running example in this paper and in the experiment. The practice of, or the advocacy for, using algorithms as decision-makers is however wider in scope. For example, see Lee & Floridi (2021) on machine learning algorithms making credit risk evaluations and mortgage lending decisions; Kleinberg et al. (2018) on machine learning predictions aiding bail decisions; Cai et al. (2019) on machine learning utilized in image retrieval systems for medical decision making.

about productivity. This is the stage at which unequal algorithmic evaluation enters: participants in the disadvantaged Blue group face a lower hiring probability because participants in the advantaged Green group are hired over a wider range of test scores due to more favorable prior beliefs. Participants are explicitly informed of this structure and know that their investment may affect their test scores, which in turn may affect hiring outcomes under the algorithmic rule. The design is intentionally stylized and transparent. Rather than reproducing the full complexity and opacity of real-world hiring algorithms, our aim is to isolate a clear form of unequal opportunity in order to study how individuals respond when investment can improve their prospects.

We find that participants in the disadvantaged Blue group invest 17.7% more than participants in the advantaged Green group, consistent with compensatory investment rather than under-investment. To shed light on the mechanisms underlying this behavior, we elicit participants’ self-reported motivation in a post-experiment questionnaire. Group identity emerges as the most frequently cited factor influencing investment decisions, exceeding considerations such as expected bonus size, investment risk, or information about predecessors’ performance. These responses are suggestive of compensatory and income-targeting reactions to disadvantage (Fongoni 2024), and indicate that participants’ awareness of the group-based evaluation rule plays an important role in shaping investment behavior.

We then test two classes of interventions designed to encourage investment. The first involves information-based encouragement through social comparison: in the “Successful Role Model” and “Unsuccessful Role Model” treatments, participants observe the investment outcomes of five in-group predecessors, motivated by a large literature documenting the effectiveness of role models in reducing bias and shaping aspirations and effort (Dasgupta & Asgari 2004, Beaman et al. 2009, Washington 2008, Beaman et al. 2012). The second intervention modifies the payoff structure: in the “No Risk” treatment, participants only incur the cost of investment if they are hired, mimicking income-contingent repayment schemes.² While the role model treatments have limited impact on investment decisions, the “No Risk” treatment increases investment among participants in the disadvantaged Blue group by 24%, suggesting that changing the economic consequences of investment may be more effective than social-comparison information in motivating self-investment under unequal algorithmic evaluation.

Our study contributes to several strands of literature. First, we contribute to the literature on statistical discrimination and behavioral responses to unequal evaluation (Arrow et al. 1973, Fang & Moro 2011, Patty & Penn 2023). In particular, the design of the hiring algorithm in the current study is most closely related to the model in Patty & Penn (2023), which formalizes how algorithms may unintentionally generate discriminatory outcomes through biased priors. Instead of focusing on hiring outcomes or normative assessments of algorithmic fairness, we examine how group-based priors shape individuals’ costly investment decisions prior to evaluation. Our results lend support to the “over-efforting” or “twice-as-hard” response to disadvantage (Caputo 2002, 2007, Evans et al. 2019, Isik et al. 2021, Gutman & Younas 2025), whereby individuals facing disadvantage increase effort to compensate. This contrasts with self-fulfilling prophecy mechanisms, in which anticipated discrimination discourages investment and reinforces initial disparities (Lundberg & Startz 1983, Coate & Loury 1993, Fryer et al. 2005, Fang & Moro 2011, Patty & Penn 2023).

²The idea of income-contingent loan schemes can be traced back to Friedman & Kuznets (1945), who proposed a novel approach to financing professional education. Friedman (1955) later expanded the concept to apply more broadly to higher education, arguing that repayments tied to future income could reduce barriers to access while aligning public and private returns on education. The income-contingent model has since been applied in a variety of policy contexts. These include drought relief for farmers (Botterill et al. 2017), where repayments are tied to future income or crop yields; fines for white-collar crimes (Chapman & Denniss 2006), which adjust penalties based on offenders’ income; and microfinance schemes for low-income households (Gans & King 2006), where loan repayment schedules are linked to the borrower’s capacity to pay.

Second, this study contributes to a growing literature on how individuals perceive and respond to algorithmic decision-makers, particularly in hiring contexts. Prior research finds that individuals often evaluate algorithmic and human decision-makers differently, with concerns about fairness, skepticism toward algorithmic screening, and preferences for human evaluators shaping attitudes toward algorithmic hiring (Dineen et al. 2004, Araujo et al. 2020, Zhang & Yencha 2022). Acceptance of algorithmic screening tools also varies systematically across individuals, with higher levels of education and income associated with greater acceptance of algorithmic hiring technologies (Zhang & Yencha 2022). Experimental evidence also suggests that discrimination by algorithms may evoke less moral outrage than comparable discrimination by humans (Bigman et al. 2023), and that transparency alone does not necessarily reduce algorithm aversion (Dargnies et al. 2024). While much of this literature focuses on attitudes or stated preferences, fewer studies examine how individuals adjust costly behavior in response to algorithmic evaluation. Notable exceptions study reliance on algorithmic advice and belief updating in judgment tasks (Burton et al. 2020, Chugunova & Sele 2022, Bayer & Renou 2026), as well as labor supply responses at the application stage (Avery et al. 2024). We extend this work by shifting attention to a different behavioral margin: costly, ex-ante investment decisions made prior to evaluation. In our incentivized experiment, participants cannot opt out of algorithmic evaluation and must decide how much to invest in improving their prospects in the presence of a known group-based evaluation rule.

Third, we contribute to the literature on interventions aimed at mitigating the behavioral consequences of biased or unequal evaluation. Existing work highlights the potential of financial instruments such as subsidies, guarantees, and income-contingent loans to counteract disincentives created by discriminatory environments (Chapman 2006, Dianat et al. 2022). Consistent with this perspective, our “No Risk” treatment, which conditions repayment on hiring and therefore reduces the downside exposure associated with investing, substantially increases investment among participants in the disadvantaged Blue group. In contrast, information-based encouragement through in-group role models has limited impact. This stands in tension with a large literature documenting strong role model effects in human-mediated contexts, including reducing bias (Dasgupta & Asgari 2004, Beaman et al. 2009), raising educational aspirations and attainment (Beaman et al. 2012, Cheryan et al. 2009), shaping political behavior (Washington 2008), and reducing crime (Iyer et al. 2012). A strand of the literature investigates the effect of “negative role models,” emphasizing that role model effectiveness depends critically on contextual fit and perceived relevance (Lockwood et al. 2002, ul Amin et al. 2025). Drawing on post-experiment questionnaires, we find that participants primarily cite group identity and financial considerations, rather than predecessor information, as drivers of their investment decisions. Together with the treatment effects, this suggests that in transparent unequal evaluation environments, social-comparison interventions may have limited influence when the decision rule is already known, while direct changes to the payoff consequences of investment may be more behaviorally relevant.

The remainder of the paper proceeds as follows. Section 2 describes the experimental design. Section 3 presents the main results. Section 4 concludes.

2 Experimental Design

The objective of this experiment is twofold. First, we test whether participants randomly assigned to a disadvantaged group invest differently than those assigned to an advantaged group under statistical discrimination from a hiring algorithm. Second, we test the effects of two classes of interventions on investment decisions: information-based encouragement through “Role Model” treatments and financial risk reduction through a “No Risk” treatment. Below we start by describing the information, choices, and payoffs faced by participants and the hiring

algorithm in the Control treatment. We then discuss the design of the “Role Model” and “No Risk” treatments, followed by the details of the implementation.

Control treatment: Does random group assignment affect investment?

The Participant’s Information. In the experiment, a participant i is randomly assigned to either the Blue or Green identity group, $d_i \in \{b, g\}$. Participants are informed that the hiring algorithm applies different prior beliefs about each group’s productivity as shown in Table 1. Let $y_i \in \{0, 1\}$ denote a participant’s productivity, where $y_i = 1$ indicates that a participant is a top performer and $y_i = 0$ indicates that a participant is an average performer. The algorithm assigns a 50% prior probability that participants from the Green group are top performers, while participants from the Blue group are assigned a prior probability of 20%.³

Independent of the algorithm’s prior beliefs, participants are informed that true productivity is randomly assigned: ex ante, each participant has a 50% chance of becoming a top performer and a 50% chance of becoming an average performer.

Participants do not directly observe their productivity type. They are informed that a pre-employment test generates an informative but noisy signal of productivity. As described in Table 2, a participant’s test score x_i takes values in $\{A, B, C\}$, where A is the highest score and C the lowest. Conditional on being an average performer, a participant has a 70% chance of receiving the lowest score C and a 30% chance of receiving the intermediate score B. Conditional on being a top performer, a participant has a 70% chance of receiving the highest score A and a 30% chance of receiving the intermediate score B.

The Participant’s Choice. Participants begin with an initial endowment and decide how much to invest, I_i , to improve their hiring prospects. Each additional unit of investment increases the probability that the participant becomes a top performer by 1.5 percentage points, thereby increasing the probability of receiving a higher test score. Investment is costly and is deducted from the participant’s endowment. If hired by the algorithm, the participant receives a positive payoff ω ; otherwise, they receive zero.

The Hiring Algorithm’s Information and Choice. The hiring algorithm acts as the potential employer and decides whether to hire each participant. It receives net payoff $\pi > 0$ from hiring a top performer, payoff $-\omega < 0$ from hiring an average performer, and zero from not hiring. We exogenously impose that the net benefit of hiring a top performer (π) is 1.5 times the wage (ω) paid to hire a participant. The hiring algorithm’s objective is therefore to maximize expected payoffs by hiring sufficiently likely top performers.

Formally, the algorithm computes the posterior probability that a participant is a top performer conditional on group identity and test score, and hires if this probability exceeds a fixed threshold. The derivation of this threshold is provided in Appendix A.1.

The algorithm’s objective implies a simple group-specific hiring rule. Participants with a test score of A are always hired regardless of group identity, while those with a score of C are never hired. For the intermediate score B, hiring depends on group membership: Green participants are hired whereas Blue participants are not, reflecting the algorithm’s less favorable prior beliefs about the Blue group. Table 3 summarizes the resulting hiring decisions by group and test score.

³The algorithm’s group-based priors are intended to capture a stylized form of statistical discrimination rather than taste-based discrimination. In our setting, the algorithm has no intrinsic preference over groups and seeks only to maximize expected hiring payoffs; group identity enters solely through beliefs about productivity. These beliefs are exogenously imposed to isolate behavioral responses to unequal algorithmic evaluation, abstracting from how such priors arise in practice (e.g., through biased training data or historical disparities).

Table 1: The hiring algorithm’s belief about group productivity rates

	$d_i = b$	$d_i = g$
$y_i = 1$	0.2	0.5
$y_i = 0$	0.8	0.5

Note: $d_i \in \{b, g\}$ indicates whether a participant is in the Blue (b) or Green (g) identity group. y_i denotes a participant’s productivity, with $y_i = 1$ denoting that a participant is a top performer, $y_i = 0$ denoting that a participant is an average performer.

Table 2: The group-blind testing technology

	$y_i = 1$	$y_i = 0$
$P(x_i = A y_i)$	0.7	0
$P(x_i = B y_i)$	0.3	0.3
$P(x_i = C y_i)$	0	0.7

Note: x_i denotes a participant’s test score which takes values in $\{A, B, C\}$, with A being the highest score and C the lowest. $P(x|y_i)$ denotes the probability that a participant receives a certain test score conditional on their true productivity. For instance, conditional on a participant being a top performer, the probability of them receiving a score of A is 70%, a score of B is 30%, and a score of C is 0%. Conditional on a participant being an average performer, the probability of them receiving a score of A is 0%, a score of B is 30%, and a score of C is 70%.

After the algorithm makes its hiring decision, participants are informed of the outcome and their corresponding payoff.

Sequence. The sequence of events in the Control treatment is summarized below.

1. Each participant i is randomly assigned to either the Blue or Green identity group, $d_i \in \{b, g\}$.
2. Participants are informed of the hiring algorithm’s group-based prior beliefs, the testing technology, and the algorithm’s hiring rule. Participants then decide how much to invest in improving their hiring prospects using the initial endowment.
3. After the participant’s investment decision, their test score $x_i \in \{A, B, C\}$ is realized according to Table 2, where $x_i = A$ is the highest score and $x_i = C$ is the lowest.
4. Observing the participant’s group d_i and test score x_i , the hiring algorithm decides whether to hire the participant according to the hiring rule shown in Table 3. The participant is then informed of the hiring decision and receives their corresponding payoff.

Table 3: Algorithmic hiring decisions by group and test score

	$x_i = A$	$x_i = B$	$x_i = C$
$d_i = g$	Hire	Hire	Not hire
$d_i = b$	Hire	Not hire	Not hire

Note: x_i denotes a participant’s test score, which takes values in $\{A, B, C\}$, with A being the highest score and C the lowest. Hiring decisions follow from the algorithm’s posterior beliefs about productivity conditional on group identity and test score, using a hiring threshold of 40% (see Appendix A.1). Participants with score A are always hired, while those with score C are never hired. For the intermediate score B, Green participants are hired but Blue participants are not, reflecting the algorithm’s less favorable prior beliefs about the Blue group.

Treatment arms: How can participants be encouraged to invest?

The experimental procedure is identical to that in the Control treatment except for the following variations.

“Role Model” treatments: Prior to making their investment decisions, participants are shown the investment decisions and outcomes of five pilot-study participants from their own group (Blue/Green). Participants are randomly assigned to one of the following two information conditions: the “Successful Role Model” treatment or the “Unsuccessful Role Model” treatment. In the “Successful Role Model” treatment, participants observe previous in-group participants who invested positive amounts and were hired. In the “Unsuccessful Role Model” treatment, participants observe previous in-group participants who invested positive amounts but were not hired. Figure A2 presents screenshots of the four possible information screens shown to participants. The decisions and outcomes displayed in these screens are based on actual pilot-study participants.⁴

“No Risk” treatment: Prior to making investment decisions, participants are informed that their chosen investment cost is deducted from their endowment only if they are hired, in contrast to the Control treatment where investment costs are deducted regardless of hiring outcome. Importantly, this treatment changes both the downside risk associated with investment and the expected payoff from investing. Figure A3 presents screenshots of the participants’ decision screens with the “No Risk” reminder.

Implementation

The study was conducted in August 2025 on Prolific. Participants were recruited to approximate national representativeness of the U.S. population across age, gender and ethnicity.⁵ It consists of a practice round, two formal (incentivized) rounds, followed by a demographic questionnaire. Figure 1 depicts the treatment assignment of the experiment. As illustrated in Figure 1, the experiment has both between- and within-subject components. With an estimated effect size of 0.2,⁶ and with $\alpha = 0.05$, power = 0.80, the power analysis predicts that having 200 observations in each treatment would provide sufficient analytical power.

To achieve the required sample size, under the constraint that we cannot undo an information treatment, we randomly assign each participant to one of the two group identities: Blue or Green, and one of the following four conditions: (1) first round: Control, second round: random assignment to “No Risk”, “Successful Role Model” or “Unsuccessful Role Model”; (2) first round: “No Risk”, second round: Control; (3) first round: “Successful Role Model”, second round: “Unsuccessful Role Model”; (4) first round: “Unsuccessful Role Model”, second round: “Successful Role Model.”⁷ Across the two decision rounds, we have 275 observations in the Control treatment; 182 in the “No Risk” treatment; 325 in the “Successful Role Model”

⁴The role model examples were selected from pilot-study participants within the same identity group who invested positive amounts. To improve comparability across the “Successful Role Model” and “Unsuccessful Role Model” treatments, the selected examples were chosen to display broadly similar investment levels across the two conditions within each group. Participants were informed that the examples came from previous participants in their own group, but were not informed about the selection procedure.

⁵Prolific constructs representative samples by matching participant subgroups to U.S. Census Bureau population shares. See <https://researcher-help.prolific.com/en/article/95c345>, accessed 02-17-2026.

⁶The assumed effect size reflects a conservative benchmark, chosen in the absence of closely comparable experimental estimates in algorithmic hiring settings, and is intended to ensure sensitivity to modest treatment effects.

⁷This design also ensures that a participant would never experience both the “No Risk” and either one of the “Role Model” treatments to avoid confounding the treatment effects.

treatment and 323 in the “Unsuccessful Role Model” treatment.⁸ Participants first learn about the relevant information. They then go through a series of attention checks to ensure that they understand the sequence of events, hiring rule and payoff calculations.⁹ Afterwards, they undergo a practice round followed by two formal investment rounds.

The survey is concluded with a short demographic questionnaire. We collect demographic variables which include age, gender, education, employment status, income, risk attitude, trust in major technology companies, and whether the participant had interacted with AI or algorithms that served as decision-makers. The full survey is appended in Appendix B.




Figure 1: Treatment assignment

The payoff parameters are set as follows. 100 ECUs (Experimental Currency Units) equate 1 USD. At the start of the experiment, each participant is endowed with 20 ECUs to invest in improving their probabilities of becoming a top performer. As described, for every 1 ECU invested, the probability the participant would become a top performer increases by 1.5 percentage points. A participant who is hired by the algorithm receives a wage (ω) of 100 ECUs. A participant’s bonus payment from a round is determined by the endowment minus their investment plus their wage if hired. In the “No Risk” treatment, a participant’s investment is

⁸A potential concern with the two-round design is that participants may adjust their behavior after observing the hiring outcome from the first formal round. We adopted the mixed within- and between-subject design to increase statistical power while maintaining the constraint that information treatments cannot be “undone” once experienced. To account for potential dependence across observations from the same participant, all regression specifications cluster standard errors at the participant level and include round fixed effects. In addition, we conduct a pre-registered robustness analysis using only first-round observations, where no prior hiring outcome has been observed.

⁹In the “Role Model” treatments, participants go through additional attention checks answering questions on the previous participants’ investment decisions and outcomes before proceeding to making their own investment decisions.

only deducted from their endowment if they’re hired. We randomly select the hiring outcome from one of the two formal rounds to determine the participants’ actual bonus payments. The median completion time of the experiment is 9 minutes 21 seconds, each participant receives a participation fee and bonus payment of 2.37 USD on average.

Hypotheses

The purpose of our hypotheses is twofold: (1) to examine whether participants facing disadvantaged algorithmic evaluation invest differently from advantaged participants, and (2) to examine the effects of each treatment on investment decisions across groups. We pre-registered all hypotheses in the AsPredicted registry under ID 243245 [<https://aspredicted.org/f8jk-5xns.pdf>].

With Hypothesis 1, we examine whether participants assigned to the disadvantaged Blue group invest differently in the presence of unequal algorithmic evaluation. Under the experimental payoff structure, expected earnings are maximized by investing fully for participants in both groups regardless of group assignment (see Appendix A.2). Standard expected-payoff maximization would therefore predict no difference in investment across groups. However, the existing literature offers competing behavioral predictions. On the one hand, anticipated discrimination may discourage investment and generate self-fulfilling prophecies (Glover et al. 2017). On the other hand, disadvantaged individuals may respond through compensatory or “twice-as-hard” effort (Caputo 2002). Given these competing mechanisms, the direction of the group effect is theoretically ambiguous *ex ante*.

Hypothesis 1 (Group Effect in Control): *In the Control treatment, participants in the disadvantaged Blue group invest differently from participants in the advantaged Green group.*

With Hypothesis 2, we test the effects of the “Role Model” treatments. Existing literature suggests that observing successful in-group peers may shape aspirations and effort decisions (Beaman et al. 2012, Cheryan et al. 2009), while observing unsuccessful peers may reduce motivation (Lockwood et al. 2002, ul Amin et al. 2025). In the context of the present experiment, “successful” role models refer to previous participants who invested positive amounts and were hired by the algorithm, while “unsuccessful” role models refer to participants who invested positive amounts but were not hired. Since the hiring rule is already fully disclosed, these treatments are intended to test whether social-comparison information affects investment behavior beyond the direct economic incentives embedded in the evaluation rule.

Hypothesis 2a (“Successful Role Model” Treatment Effect): *Within each group (Blue or Green), participants in the “Successful Role Model” treatment will invest more than those in the Control treatment.*

Hypothesis 2b (“Unsuccessful Role Model” Treatment Effect): *Within each group (Blue or Green), participants in the “Unsuccessful Role Model” treatment will invest less than those in the Control treatment.*

With Hypotheses 3 and 4, we test the effect of the “No Risk” treatment and compare it with the Control and “Role Model” treatments. Because the “No Risk” treatment reduces downside exposure associated with investing and simultaneously increases the expected return to investment, we posit that participants would invest more in the “No Risk” treatment than in the Control treatment. Since the “Role Model” treatments alter social-comparison information but do not directly change the payoff structure, we additionally hypothesize that investment

will be higher in the “No Risk” treatment than in the “Role Model” treatments.

Hypothesis 3 (“No Risk” Treatment Effect): *Within each group (Blue or Green), participants in the “No Risk” treatment will invest more than those in the Control treatment.*

Hypothesis 4 (“No Risk” vs. “Role Model”): *Within each group (Blue or Green), participants in the “No Risk” treatment will invest more than those in either “Role Model” treatment.*

Finally, with Hypothesis 5, we conduct a heterogeneity analysis to examine whether treatment effects differ across groups. Since both the “Role Model” and “No Risk” treatments are motivated by interventions designed to mitigate the behavioral consequences of disadvantage, we hypothesize that these treatments will have stronger effects among participants in the disadvantaged Blue group.¹⁰

Hypothesis 5 (Heterogeneity Analysis): *The effects of the “Role Model” and “No Risk” treatments differ across the two groups, with stronger treatment effects among the disadvantaged Blue group.*

3 Results

Summary Statistics

Table 4 presents summary statistics. Participants are, on average, 45.8 years old and roughly split between male and female. Approximately 54.5% hold at least a Bachelor’s degree, and 69.3% are employed, with the median income falling within the bracket “Greater than or equal to 50,000 USD and less than 75,000 USD.” Participants scored an average of 6.5 on the 0–10 risk attitude scale, and 5.7 on the 0–10 trust scale in response to the question: “In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way?” When asked, “Have you ever applied for a job or other opportunity where you believe your application was screened or reviewed by an AI system or algorithm?” around 35.8% of participants selected the response: “Yes, I have submitted an application where I believe my application was screened by an AI system or algorithm.” Out of the participants who selected yes, replying to the follow-up question: “If so, in which contexts have you interacted with the AI system or algorithm? Select all that apply.” 96% reported that they had interacted with an AI system or algorithm during job applications or through hiring platforms; 35.4% reported having had such interaction in the context of credit scoring or loan applications. We report the summary statistics by identity group (Blue vs. Green), and the summary statistics by treatment (Control, Successful Role Model, Unsuccessful Role Model, and No Risk) respectively in Table A1 and Table A2 in the Appendix. Across both tables, we observe no systematic differences in demographic characteristics, prior experience with AI or algorithm, trust in major tech companies, or risk attitudes.

Figure 2 shows the distribution of investment decisions by group across treatment conditions. Across the four treatments, investments for both groups cluster around 0, 10, and 20 units, with Blue participants more likely to invest the maximum amount of 20. While the distribution of Green participants’ investments remains relatively stable across treatment conditions, the distribution of Blue participants’ investments exhibits a sharp spike at the maximum investment level in the “No Risk” treatment.

¹⁰Additionally, proof that the disadvantaged Blue group faces greater payoff variance under the experimental environment is provided in Appendix A.2.

Table 4: Summary statistics

Variable	Summary (Mean, SD)
Age	45.79 (15.62)
Risk attitude	6.50 (2.55)
Trust in major tech companies	5.72 (2.40)
Male	49.4%
Employed	69.3%
Prior experience with AI or algorithm	35.8%
Education: Some high school or less	1.1%
Education: High school diploma or GED	13.4%
Education: Some college, but no degree	16.6%
Education: Associates or technical degree	14.3%
Education: Bachelor’s degree	35.7%
Education: Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, etc.)	18.8%
Income: Less than 25k	24.2%
Income: 25-50k	25.5%
Income: 50-75k	19.5%
Income: 75-100k	14.6%
Income: 100-125k	4.9%
Income: Greater than or equal to 125k	11.4%
Observations	553

Note: Age is measured in years as a continuous variable. Risk attitude is recorded on a 0–10 scale, with higher scores indicating greater willingness to take risks. Trust in major tech companies is based on responses to the question: “In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way?”, also on a 0–10 scale. Male, Employed, and Prior experience with AI or algorithm are binary variables coded as 1 if the participant identifies as male, reports being employed (full-time or part-time), or has previously submitted an application believed to be screened by an AI system or algorithm.

Effect of Group Advantage on Investment

To test Hypothesis 1, we estimate the following OLS regression using data from the Control treatment:

$$Invest_i = \beta_0 + \beta_1 Blue_i + X_i' \gamma + \varepsilon_i$$

where $Invest_i$ is the outcome variable that denotes the investment amount of participant i ; β_0 is the investment level of the Green group; $Blue_i$ is a dummy variable that equals 1 if the participant is in the Blue group, and 0 if the participant is in the Green group. X_i' is a vector of pre-registered controls including round fixed effects, demographics, risk preferences, trust in major tech companies and prior experience with AI or algorithm. Standard errors are clustered at the individual level. Table 5 shows the regression results. In Column 1, we report our main specification controlling only for round fixed effects, in Column 2, we include demographic controls, and in Column 3, risk attitude, prior experience with AI or algorithm and trust in major technology companies are also controlled for.

Across all three specifications, we observe a statistically significant, positive effect of being in the Blue group on participants’ investment amount. Our main specification in Column (1) reports that Green participants invest 12.4 ECUs out of 20 on average, and that compared to Green participants, Blue participants invest 2.2 ECUs more. In other words, participants in the disadvantaged Blue group invest around 17.7% more than their counterparts in the advantaged Green group.

This positive effect is also robust under the pre-registered alternative specification where we exclude participants with completion time outside the 10th and 90th percentile. In another




Figure 2: Investment distribution by group

pre-registered robustness specification, we restrict the analysis to first-round observations only. The estimated coefficient remains positive and similar in magnitude to our main specification, although it is no longer statistically significant, likely reflecting the substantially smaller sample size available in the first-round analysis (140 observations in the Control treatment). The regression results of the alternative specifications are reported in Table A4 and Table A6 in the Appendix. One additional pre-registered exclusion criterion involves removing participants who failed attention checks. However, due to programming constraints on the survey platform, failed attempts were not recorded once participants eventually provided correct responses. As a result, this exclusion cannot be implemented ex post, and all participants who successfully completed the attention checks are retained in the analysis.

Taken together, the results suggest that participants do not disengage when assigned to the disadvantaged group. Instead, they increase investment, consistent with a compensatory response to unequal evaluation.

Result 1 (support for H1): *In the Control treatment, participants in the disadvantaged Blue group invest more than participants in the advantaged Green group.*

“Role Model” & “No Risk” Treatment Effects

To test Hypotheses 2, 3, 4 & 5, we estimate the following OLS regression:

$$\begin{aligned} Invest_i = & \beta_0 + \beta_1 SRM_i + \beta_2 URM_i + \beta_3 NR_i + \beta_4 Blue_i + \beta_5 (SRM_i \times Blue_i) \\ & + \beta_6 (URM_i \times Blue_i) + \beta_7 (NR_i \times Blue_i) + X_i' \gamma + \varepsilon_i \end{aligned}$$

where $SRM_i = 1$ if participant i is in the “Successful Role Model” treatment, $URM_i = 1$ if participant i is in the “Unsuccessful Role Model” treatment, $NR_i = 1$ if participant i is in the “No Risk” treatment, $Blue_i = 1$ if participant i is in the Blue group, and 0 otherwise. X_i' is a vector of pre-registered controls including round fixed effects, demographics, risk preferences, trust in major tech companies and prior experience with AI or algorithm. Standard errors

Table 5: Effect of group (Blue vs. Green) on investment in control treatment

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group (vs Green)	2.199** (0.928)	2.409** (0.984)	2.184** (1.015)
Constant	12.415*** (0.820)	7.961*** (2.668)	8.458** (3.788)
Round Control	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	276	276	276
R ²	0.041	0.136	0.235

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

are clustered at the individual level. Our coefficients of interest are β_1 and $\beta_1 + \beta_5$ (H2a), β_2 and $\beta_2 + \beta_6$ (H2b), β_3 and $\beta_3 + \beta_7$ (H3), which yield the treatment effects within the Green group and the Blue group respectively and allow us to compare effects across treatments (H4). The coefficients of the interaction terms, $\beta_5, \beta_6, \beta_7$, test heterogeneity in treatment effects across groups (H5). Table 6 reports the results from our main specification and the linear combination tests.¹¹

From Table 6, we first note that both “Role Model” treatments have no statistically significant effect on investment for either group in the main specification. This finding is consistent with the robustness check reported in Table A5 (Appendix), where we exclude participants whose completion time falls outside the 10th to 90th percentile. However, when we restrict the analysis to first-round observations only (Table A7, Appendix), we observe a negative effect of the “Unsuccessful Role Model” treatment on investment for both groups, with participants investing approximately 1.9 ECUs less than those in the Control treatment. By contrast, the “Successful Role Model” treatment remains ineffective across specifications.

The first-round specification is informative because it eliminates any possibility that participants adjust their behavior after observing the outcome of a previous round. Comparing the first-round and pooled specifications suggests that exposure to unsuccessful in-group predecessors may initially influence investment decisions. However, this pattern is not robust across specifications and is not present in the full sample. We therefore interpret that the role model treatments do not generate consistent effects on investment behavior.

Result 2 (no support for H2a, H2b): *Within each group (Blue or Green), participants in the “Role Model” treatments do not invest differently than those in the Control treatment in our main specification.*

We next turn our attention to the effect of the “No Risk” treatment. As discussed above, this treatment changes both the downside risk associated with investment and the expected payoff from investing, and should therefore be interpreted as reducing the downside exposure associated with investment rather than as a pure risk treatment. In our main specification, being in the “No Risk” treatment does not significantly affect Green participants’ investment decisions. However, for participants in the disadvantaged Blue group, the effect is positive and significant. Compared to Blue participants in the Control treatment, being in the “No Risk” treatment increases a Blue participant’s investment by approximately 3.5 ECUs (24%). Compared to Green participants who are also in the “No Risk” treatment, Blue participants in the “No Risk” treatment invest 2.3 ECUs (17%) more.

This positive, heterogeneous effect is robust under both of our alternative specifications where we exclude participants with completion time outside the 10th and 90th percentile (Table A5) and use data only from the first decision round (Table A7). In the alternative specification where we exclude participants whose completion time falls outside the 10th to 90th percentile (Table A5), we also find a positive and significant effect of the “No Risk” treatment on Green participants’ investment amount. However, such effect is not detected in the other specifications. Therefore, while the strongest and most robust treatment response is observed among Blue participants, the results should not be interpreted as implying that Green participants are entirely unresponsive to changes in the payoff structure.

Result 3 (partial support for H3): *The “No Risk” treatment has a positive and robust*

¹¹Table A3 (Appendix) reports treatment effects under all three pre-registered specifications. In Column 1, we report our main specification controlling only for round fixed effects, in Column 2, we include demographic controls, and in Column 3, risk attitude, prior experience with AI or algorithm and trust in major technology companies are also controlled for.

effect on Blue participants' investment amounts. The effect on Green participants is less robust across specifications.

Among Green participants, neither the "Successful Role Model" nor the "Unsuccessful Role Model" treatment significantly affects investment relative to the Control treatment in the main specification, and the "No Risk" treatment also does not yield a statistically significant change. In contrast, among Blue participants, the interaction term for Blue $\times$ No Risk is positive and statistically significant. The linear combination test indicates that Blue participants in the "No Risk" treatment invest significantly more than Blue participants in the Control condition.

To directly compare the "No Risk" treatment with the two "Role Model" treatments within the Blue group, we conduct linear combination tests of the relevant coefficients. These tests show that investment under "No Risk" is significantly higher than under either role model treatment for Blue participants. No such differences are observed within the Green group. These results are consistent with the interpretation that, in this transparent unequal evaluation environment, direct changes to the economic consequences of investment are more behaviorally relevant than social-comparison information.

Result 4 (partial support for H4): *Blue participants in the "No Risk" treatment invest more than those in either "Role Model" treatment. There is no difference in Green participants' investment amount across "No Risk" and "Role Model" treatments.*

Result 5 (partial support for H5): *The "Role Model" treatments have no significant effect in either group in the main specification. The effect of the "No Risk" treatment differs across groups and is strongest among participants in the disadvantaged Blue group.*

Table 6: Treatment effects by group and linear combination tests

	Regression coefficients		Linear combination tests	
	Estimate	Std. Error	Estimate	Std. Error
Main effects				
Blue Group	2.317**	(0.929)		
Successful RM (Green)	0.503	(0.847)	0.503	(0.856)
Unsuccessful RM (Green)	0.152	(0.889)	0.152	(0.856)
No Risk (Green)	1.216	(0.750)	1.216	(0.997)
Interactions				
Blue $\times$ Successful RM	0.551	(1.145)	1.055	(0.864)
Blue $\times$ Unsuccessful RM	-0.672	(1.232)	-0.520	(0.866)
Blue $\times$ No Risk	2.264**	(1.046)	3.480***	(1.022)
Constant	11.744***	(0.678)		
Observations			1,106	
R ²			0.052	

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. The first two columns report OLS regression coefficients from Column (1) of Table A3. The final two columns show linear combination tests comparing each treatment condition to the Control group within Blue and Green participants. Standard errors are clustered at the participant level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

For completeness, Appendix Table A16 reports hiring rates by group and treatment. As expected given the algorithm's group-specific hiring rule, hiring rates remain higher among Green participants across all treatments, despite the higher investment levels observed among

Blue participants.

Exploratory Analyses

To better understand individual-level heterogeneity in investment decisions, we explore whether demographic characteristics and psychological traits moderate the effects of unequal algorithmic evaluation and the experimental treatments. We estimate a series of models that include interactions between treatment assignments and participants’ age, gender, education, employment status, income, prior experience with algorithms, risk attitudes, and trust in technology companies. Regression results of these analyses are reported in Table A8 to Table A15 in the Appendix. These analyses are exploratory and should be interpreted as suggestive rather than confirmatory.

We find that gender matters: participants identifying as male invest more than the others (Gender coefficient: 1.543, $p < 0.05$), even after controlling for other demographic and psychological variables. Employment status also plays a role. Participants who are currently employed (either full-time or part-time) tend to invest more overall than those who are not (Employment coefficient: 3.740, $p < 0.01$). This positive association is weaker in the “No Risk” treatment (Employment $\times$ No Risk: -2.163 , $p < 0.1$), though this interaction effect is only marginally significant and should be interpreted cautiously.

Similarly, prior experience with algorithms as decision-makers is positively associated with investment (coefficient: 2.206, $p < 0.01$). Participants with higher education appear to invest marginally more in response to the “Successful Role Model” treatment (Education $\times$ Successful RM: 1.954, $p < 0.1$), while participants with higher trust in major technology companies invest marginally less under both the “Successful Role Model” (Trust $\times$ Successful RM: 4.026, $p < 0.1$) and “No Risk” (Trust $\times$ No Risk: 3.711, $p < 0.1$) treatments.

By contrast, we do not find significant moderating effects for income or general risk attitude, with the exception of a marginally negative interaction between income and the “Successful Role Model” treatment (Income $\times$ Successful RM: -0.465 , $p < 0.1$). It is also notable that risk attitude neither predicts investment nor significantly interacts with the “No Risk” treatment. Since the “No Risk” treatment changes both downside risk and expected payoffs, this result suggests that the treatment effect does not appear to be driven primarily by baseline risk preferences. Instead, participants may be responding more broadly to the reduced downside exposure and improved payoff consequences of investing.


To further probe the mechanisms behind investment decisions, in the post-experiment questionnaire, we ask participants to indicate which factors motivated their choices. Participants can select all that apply out of the following options presented in random order: “My group identity (whether I was assigned to the Green or Blue group),” “Whether I have to pay for the investment cost when I’m not hired,” “How much previous participants invested,” “Whether previous participants were hired,” “The starting amount I had available to invest (20 ECUs),” “The potential bonus I could get if hired (100 ECUs),” “Other (please specify in the textbox below).” Options “How much previous participants invested,” and “Whether previous participants were hired” were only displayed to participants who underwent at least one of the “Role Model” treatments. The participants’ responses separated by group are shown in Figure 3. The most frequently cited factor is group identity (whether the participant is assigned to the Blue or Green group), rather than expected bonus size, investment cost, or information about predecessors’ performance or outcome. This suggests that the group-based evaluation rule was highly salient to participants’ investment decisions.

More specifically, this pattern is consistent with compensatory investment responses to disadvantage, whereby individuals facing less favorable evaluation conditions increase investment in an attempt to offset disadvantage. In this sense, group identity may serve as a reference

point for expected outcomes, motivating higher investment among disadvantaged participants in line with income-targeting or gap-closing mechanisms discussed in the inequality and effort literature (e.g., Fongoni 2024). However, because these responses are self-reported and elicited after the investment decisions, they should be interpreted as descriptive evidence rather than direct causal evidence on mechanisms.

Open-ended responses from participants who selected “Other (please specify in the textbox below)” in the motivation question (Figure 3; $N = 10$) provide additional descriptive insight into participants’ reasoning. Given the small size of this subsample, these patterns should be interpreted with caution. We used ChatGPT to group responses into broad thematic categories based on their content, without pre-specifying themes. This process yielded three recurring categories: (1) cost-benefit reasoning, such as “statistics and probability of maximizing my profit” and “risk only 20 cents to potentially gain 1 dollar”; (2) goal-oriented motivations, such as “I wanted to get hired” and “doing whatever it took to increase my likelihood of being hired”; and (3) pessimism and algorithmic aversion, including statements like “no motivation knowing it was a rigged game” and “the AI was designed to not pay out.” These responses again suggest that some participants interpreted the environment through the lens of incentives, while others perceived the evaluation rule as unfair or discouraging.

Figure 3: Self-reported motivation for investment



Note: The exact wording of the question to elicit participants’ motivation for investment is: “In the experiment, what factor(s) below motivated your investment choice? Select all that apply.” The choices presented in random order that participants can select from are: “My group identity (whether I was assigned to the Green or Blue group),” “Whether I have to pay for the investment cost when I’m not hired,” “How much previous participants invested,” “Whether previous participants were hired,” “The starting amount I had available to invest (20 ECUs),” “The potential bonus I could get if hired (100 ECUs),” “Other (please specify in the textbox below).” Options “How much previous participants invested,” and “Whether previous participants were hired” were only displayed to participants who underwent at least one of the “Role Model” treatments.

4 Conclusion

The results from our experiment reveal that participants assigned to the disadvantaged Blue group invest more than participants assigned to the advantaged Green group. This finding is consistent with the literature documenting the “over-efforting” or “twice-as-hard” effect, and suggests that unequal evaluation does not necessarily lead to disengagement when individuals can respond through costly self-investment. We also find that reducing the downside risk exposure associated with investing through the “No Risk” treatment leads to higher investment among disadvantaged participants. In contrast, the “Role Model” treatments, designed to influence behavior through social comparison, do not shift investment decisions meaningfully in the main specification.

Our findings suggest an interpretation in which behavioral responses to unequal evaluation depend not only on disadvantage itself, but also on features of the institutional environment. In our setting, participants are explicitly informed of the algorithm’s group-based hiring rule and know how investment affects their probability of receiving a favorable signal. This environment allows us to isolate how individuals respond to unequal opportunity when the decision rule is fixed and self investment remains consequential. The results suggest that disadvantaged individuals may respond to such environments through compensatory investment rather than under-investment.

The limited effect of the “Role Model” treatments is informative in this context. Since the algorithm’s hiring rule is fixed and known, participants may have interpreted predecessors’ outcomes as less relevant for their own prospects. Relative to settings in which social comparison or exposure to successful in-group predecessors affects behavior through aspirations, beliefs, or expectations about discretionary evaluators (Washington 2008, Beaman et al. 2009), such information may be less influential when the evaluator follows a pre-specified rule. This interpretation is consistent with the post-experiment questionnaire, where participants primarily cite group identity and financial considerations rather than predecessor information as factors motivating their investment decisions.

By contrast, the strong response to the “No Risk” treatment indicates that interventions which directly modify the payoff structure remain effective under unequal algorithmic evaluation. Importantly, this treatment changes both the downside risk associated with investment and the expected payoff from investing. We therefore interpret the result as evidence that disadvantaged participants are responsive to reduced downside exposure, rather than as evidence for a pure risk-preference mechanism. This interpretation is further supported by the exploratory analysis showing that baseline risk attitudes do not significantly predict investment or moderate the effect of the “No Risk” treatment. More broadly, the result suggests that when individuals face unequal evaluation rules, reducing the financial consequences of failure may be more effective than social-comparison information in encouraging self-investment.

At the same time, the higher investment levels observed among disadvantaged participants do not fully eliminate the disparities generated by unequal algorithmic evaluation. As reported in Appendix Table A16, participants in the advantaged Green group continue to experience higher hiring rates than participants in the disadvantaged Blue group across all treatment conditions. This pattern suggests that compensatory investment may mitigate, but does not fully offset, disadvantages embedded in the evaluation process.

The experiment necessarily abstracts from several features of real-world hiring processes. First, the hiring algorithm used in this study applies a fixed decision rule and does not adapt over time. In practice, many algorithmic systems update based on incoming data or operate in hybrid decision environments involving human recruiters (Tambe et al. 2019, Marabelli et al. 2021). A static design was deliberately chosen to maintain experimental control and isolate behavioral responses to a transparent form of unequal algorithmic evaluation. Allowing the

algorithm to adapt would introduce endogeneity between individual investment decisions and evolving hiring thresholds. Future work could extend this framework to dynamic settings in which algorithmic learning and individual behavior co-evolve.

Second, our study abstracts from competitive dynamics among applicants. All participants who meet a predefined qualification threshold may be hired, eliminating strategic considerations about others’ behavior. While there are examples of algorithmic screenings being adopted in non-competitive hiring contexts,¹² and that this non-competitive structure allows for clean identification of individual responses to algorithmic bias, future research could examine how these responses change in relative performance or rank-based selection environments. Such extensions would be particularly relevant in labor markets where algorithmic screening is used to allocate scarce opportunities.

Third, our design deliberately uses a transparent algorithmic rule. In many real-world settings, individuals rarely know whether an algorithm is biased, let alone the exact form of the disadvantage they face. This transparency is a limitation for external validity, but it is also central to the identification strategy: by making the unequal evaluation rule explicit, the experiment isolates behavioral responses to perceived unequal opportunity without confounding them with uncertainty about the decision process. Future research could examine how responses differ when individuals are uncertain about whether discrimination exists, how severe it is, or whether effort can overcome it.

Finally, this study focuses on individual behavioral responses rather than institutional reform. While understanding how individuals react to algorithmic disadvantage is informative, it does not address how discriminatory algorithms should be redesigned or regulated. Indeed, the hiring-rate comparisons reported in the Appendix suggest that greater self-investment alone is insufficient to eliminate disparities when unequal evaluation remains in place. Placing the burden of adjustment on individuals therefore risks reinforcing a “fix the person, not the system” narrative (Bertrand & Duflo 2017, Kline et al. 2022). Nevertheless, individual-level responses remain an important input for policy design. Many existing interventions, such as income-contingent loans, training subsidies, or financial guarantees, operate by altering the risk and payoff structure faced by individuals. Our findings suggest that such tools may help encourage investment under unequal algorithmic evaluation, but they should be viewed as complements rather than substitutes for institutional reforms that address disparities at their source.

¹²For examples of algorithmic screening outside competitive hiring contexts, see Kleinberg et al. (2018) on bail decisions and Lee & Floridi (2021) on credit risk evaluation.

References

- Ajunwa, I. & Greene, D. (2019), Platforms at work: Automated hiring platforms and other new intermediaries in the organization of work, *in* ‘Work and labor in the digital age’, Vol. 33, Emerald Publishing Limited, pp. 61–91.
- Araujo, T., Helberger, N., Kruikemeier, S. & De Vreese, C. H. (2020), ‘In ai we trust? perceptions about automated decision-making by artificial intelligence’, *AI & society* **35**(3), 611–623.
- Arrow, K., Ashenfelter, O. & Rees, A. (1973), ‘Discrimination in labor markets’, *The Theory of Discrimination* pp. 3–33.
- Avery, M., Leibbrandt, A. & Vecchi, J. (2024), ‘Does artificial intelligence help or hurt gender diversity? evidence from two field experiments on recruitment in tech’.
- Awad, E., Balafoutas, L., Chen, L., Ip, E. & Vecchi, J. (2023), ‘Artificial intelligence and debiasing in hiring: Impact on applicant quality and gender diversity’, *Available at SSRN* .
- Barocas, S. & Selbst, A. D. (2016), ‘Big data’s disparate impact’, *Calif. L. Rev.* **104**, 671.
- Bayer, R.-C. & Renou, L. (2026), ‘Interacting with man or machine: When do humans reason better?’, *Management Science* **72**(1), 594–608.
- Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R. & Topalova, P. (2009), ‘Powerful women: does exposure reduce bias?’, *The Quarterly journal of economics* **124**(4), 1497–1540.
- Beaman, L., Duflo, E., Pande, R. & Topalova, P. (2012), ‘Female leadership raises aspirations and educational attainment for girls: A policy experiment in india’, *science* **335**(6068), 582–586.
- Bertrand, M. & Duflo, E. (2017), ‘Field experiments on discrimination’, *Handbook of economic field experiments* **1**, 309–393.
- Bigman, Y. E., Wilson, D., Arnestad, M. N., Waytz, A. & Gray, K. (2023), ‘Algorithmic discrimination causes less moral outrage than human discrimination.’, *Journal of Experimental Psychology: General* **152**(1), 4.
- Bîgu, D. & Cernea, M.-V. (2019), Algorithmic bias in current hiring practices: An ethical examination, *in* ‘13th International Management Conference (IMC) on Management Strategies for High Performance. 13th International Management Conference (IMC) on Management Strategies for High Performance, Bucharest, Romania, October’.
- Binns, R. (2018), Fairness in machine learning: Lessons from political philosophy, *in* ‘Conference on fairness, accountability and transparency’, PMLR, pp. 149–159.
- Bogen, M. & Rieke, A. (2018), ‘Help wanted: An examination of hiring algorithms, equity, and bias’, *Upturn*, December **7**.
- Botterill, L. C., Chapman, B. & Kelly, S. (2017), ‘Revisiting revenue contingent loans for drought relief: government as risk manager’, *Australian Journal of Agricultural and Resource Economics* **61**(3), 367–384.
- Burton, J. W., Stein, M.-K. & Jensen, T. B. (2020), ‘A systematic review of algorithm aversion in augmented decision making’, *Journal of behavioral decision making* **33**(2), 220–239.

- Cai, C. J., Reif, E., Hegde, N., Hipp, J., Kim, B., Smilkov, D., Wattenberg, M., Viegas, F., Corrado, G. S., Stumpe, M. C. et al. (2019), Human-centered tools for coping with imperfect algorithms during medical decision-making, *in* ‘Proceedings of the 2019 chi conference on human factors in computing systems’, pp. 1–14.
- Caputo, R. K. (2002), ‘Discrimination and human capital: A challenge to economic theory and social justice’, *J. Soc. & Soc. Welfare* **29**, 105.
- Caputo, R. K. (2007), ‘Perceived work-related discrimination by women: Implications for social justice and affirmative action’, *Journal of Policy Practice* **6**(2), 5–22.
- Chapman, B. (2006), ‘Income contingent loans for higher education: International reforms’, *Handbook of the Economics of Education* **2**, 1435–1503.
- Chapman, B. & Denniss, R. (2006), Criminal reparations: using financial incentives and income contingent fines for white-collar crimes, *in* ‘Government Managing Risk’, Routledge, pp. 176–193.
- Cheryan, S., Plaut, V. C., Davies, P. G. & Steele, C. M. (2009), ‘Ambient belonging: how stereotypical cues impact gender participation in computer science.’, *Journal of personality and social psychology* **97**(6), 1045.
- Chugunova, M. & Sele, D. (2022), ‘We and it: an interdisciplinary review of the experimental evidence on how humans interact with machines’, *Journal of Behavioral and Experimental Economics* **99**, 101897.
- Coate, S. & Loury, G. C. (1993), ‘Will affirmative-action policies eliminate negative stereotypes?’, *The American Economic Review* pp. 1220–1240.
- Dargnies, M.-P., Hakimov, R. & Kübler, D. (2024), ‘Aversion to hiring algorithms: Transparency, gender profiling, and self-confidence’, *Management Science* .
- Dasgupta, N. & Asgari, S. (2004), ‘Seeing is believing: Exposure to counterstereotypic women leaders and its effect on the malleability of automatic gender stereotyping’, *Journal of experimental social psychology* **40**(5), 642–658.
- Dianat, A., Echenique, F. & Yariv, L. (2022), ‘Statistical discrimination and affirmative action in the lab’, *Games and Economic Behavior* **132**, 41–58.
- Dineen, B. R., Noe, R. A. & Wang, C. (2004), ‘Perceived fairness of web-based applicant screening procedures: Weighing the rules of justice and the role of individual differences’, *Human Resource Management: Published in Cooperation with the School of Business Administration, the University of Michigan and in Alliance with the Society of Human Resources Management* **43**(2-3), 127–145.
- Evans, D. P., Sales, J. M., Krause, K. H. & Del Rio, C. (2019), ‘You have to be twice as good and work twice as hard: a mixed-methods study of perceptions of sexual harassment, assault and women’s leadership among female faculty at a research university in the usa’, *Global health, epidemiology and genomics* **4**, e6.
- Fang, H. & Moro, A. (2011), ‘Theories of statistical discrimination and affirmative action: A survey’, *Handbook of social economics* **1**, 133–200.
- Fehr, D., Müller, D. & Preuss, M. (2024), ‘Social mobility perceptions and inequality acceptance’, *Journal of Economic Behavior & Organization* **221**, 366–384.

- Fongoni, M. (2024), ‘Does pay inequality affect worker effort? an assessment of experimental designs and evidence’, *Journal of Economic Behavior & Organization* **220**, 697–716.
- Friedman, M. (1955), ‘The role of government in education’.
- Friedman, M. & Kuznets, S. (1945), *Income from independent professional practice*, number frie54-1, National Bureau of Economic Research.
- Fryer, R. G., Goeree, J. K. & Holt, C. A. (2005), ‘Experience-based discrimination: Classroom games’, *The Journal of Economic Education* **36**(2), 160–170.
- Gans, J. S. & King, S. P. (2006), Income contingent loans for low-income households, in ‘Government Managing Risk’, Routledge, pp. 216–236.
- Glover, D., Pallais, A. & Pariente, W. (2017), ‘Discrimination as a self-fulfilling prophecy: Evidence from french grocery stores’, *The Quarterly Journal of Economics* **132**(3), 1219–1260.
- Gutman, L. M. & Younas, F. (2025), ‘Understanding the awarding gap through the lived experiences of minority ethnic students: An intersectional approach’, *British Educational Research Journal* **51**(2), 990–1008.
- Hamilton, I. A. (2018), ‘Why it’s totally unsurprising that amazon’s recruitment ai was biased against women’, *Business Insider*. Retrieved November 11.
- Houser, K. A. (2019), ‘Can ai solve the diversity problem in the tech industry: Mitigating noise and bias in employment decision-making’, *Stan. Tech. L. Rev.* **22**, 290.
- Isik, U., Wouters, A., Verdonk, P., Croiset, G. & Kusurkar, R. A. (2021), ‘“as an ethnic minority, you just have to work twice as hard.” experiences and motivation of ethnic minority students in medical education’, *Perspectives on medical education* **10**(5), 272–278.
- Iyer, L., Mani, A., Mishra, P. & Topalova, P. (2012), ‘The power of political voice: women’s political representation and crime in india’, *American Economic Journal: Applied Economics* **4**(4), 165–193.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. & Mullainathan, S. (2018), ‘Human decisions and machine predictions’, *The quarterly journal of economics* **133**(1), 237–293.
- Kline, P., Rose, E. K. & Walters, C. R. (2022), ‘Systemic discrimination among large us employers’, *The Quarterly Journal of Economics* **137**(4), 1963–2036.
- Kordzadeh, N. & Ghasemaghaei, M. (2022), ‘Algorithmic bias: review, synthesis, and future research directions’, *European Journal of Information Systems* **31**(3), 388–409.
- Lee, M. S. A. & Floridi, L. (2021), ‘Algorithmic fairness in mortgage lending: from absolute conditions to relational trade-offs’, *Minds and Machines* **31**(1), 165–191.
- Lockwood, P., Jordan, C. H. & Kunda, Z. (2002), ‘Motivation by positive or negative role models: regulatory focus determines who will best inspire us.’, *Journal of personality and social psychology* **83**(4), 854.
- Lundberg, S. J. & Startz, R. (1983), ‘Private discrimination and social intervention in competitive labor market’, *The American economic review* **73**(3), 340–347.

- Marabelli, M., Newell, S. & Handunge, V. (2021), ‘The lifecycle of algorithmic decision-making systems: Organizational choices and ethical challenges’, *The Journal of Strategic Information Systems* **30**(3), 101683.
- McCall, L., Burk, D., Laperrière, M. & Richeson, J. A. (2017), ‘Exposure to rising inequality shapes americans’ opportunity beliefs and policy support’, *Proceedings of the National Academy of Sciences* **114**(36), 9593–9598.
- Patty, J. W. & Penn, E. M. (2023), ‘Algorithmic fairness and statistical discrimination’, *Philosophy Compass* **18**(1), e12891.
- Phelps, E. S. (1972), ‘The statistical theory of racism and sexism’, *The american economic review* **62**(4), 659–661.
- Tambe, P., Cappelli, P. & Yakubovich, V. (2019), ‘Artificial intelligence in human resources management: Challenges and a path forward’, *California Management Review* **61**(4), 15–42.
- ul Amin, N., Kiani, A., Ahmad, N. & Weng, Q. (2025), ‘Motivation by role models: Do they inspire or discourage us? investigating the link between role models and entrepreneurial career intentions’, *The International Journal of Management Education* **23**(2), 101090.
- Van Esch, P., Black, J. S. & Ferolie, J. (2019), ‘Marketing ai recruitment: The next phase in job application and selection’, *Computers in Human Behavior* **90**, 215–222.
- Vincent, J. (2018), ‘Amazon reportedly scraps internal ai recruiting tool that was biased against women’, *The Verge* **10**.
- Washington, E. L. (2008), ‘Female socialization: How daughters affect their legislator fathers’ voting on women’s issues’, *American Economic Review* **98**(1), 311–332.
- Zhang, L. & Yench, C. (2022), ‘Examining perceptions towards hiring algorithms’, *Technology in Society* **68**, 101848.

A Appendix

A.1 Algorithmic Hiring Rule Derivation

Let $\pi > 0$ denote the hiring algorithm’s net payoff from hiring a top performer, $-\omega < 0$ the payoff from hiring an average performer (or, the wage of hiring any participant), and 0 when not hiring. Observing the participant’s group (d_i) and test score (x_i), the hiring algorithm hires the participant if and only if the expected utility of hiring, $EU(hire)$, satisfies:

$$EU(hire) = P(y_i = 1|d_i, x_i) \times \pi + (1 - P(y_i = 1|d_i, x_i)) \times (-\omega) > 0$$

That is, to be hired by the algorithm, $P(y_i = 1|d_i, x_i)$ must satisfy:

$$P(y_i = 1|d_i, x_i) \geq \frac{\omega}{(\pi + \omega)}$$

In the experiment, we exogenously impose that the net benefit of hiring a top performer (π) is 1.5 times the wage (ω) paid to hire a participant. Plugging the numbers in the above function, the hiring algorithm hires a participant when, conditional on the participant’s group and test score, the probability that a participant is a top performer is no less than 40%.

When a participant receives a test score of A, with the group-blind testing technology described in Table 2, the hiring algorithm could be 100% certain that this participant is a “top performer” and would hire regardless of the participant’s group identity. Similarly for when a participant receives a test score of C, the hiring algorithm could be 100% certain that the participant is an “average performer” and would never hire regardless of the participant’s group. When a participant receives a B, however, given that it is an uninformative score (both “top performers” and “average performers” are equally likely to receive a score of B), the hiring algorithm has to refer to its biased group-based priors to make a decision. Recall from Table 1 the hiring algorithm’s belief about each group’s productivity rate, where it believes that 50% of participants from the Green group are “top performers” and that only 20% of participants from the Blue group are “top performers.” For a Green participant with a score of B, the hiring algorithm rely on its priors and believes that the participant has a 50% chance of being a “top performer,” and would hire the participant given that the probability is higher than the hiring threshold of 40%. However, for a Blue participant with a score of B, given that the hiring algorithm’s prior of Blue participants being “top performers” is 20%, they would not be hired given that the probability is lower than the hiring threshold.

A.2 Marginal Payoff and Variance Derivation

With the testing technology and the hiring threshold described in Section 2 and Appendix A.1, and denote $p_0 \in \{p_g, p_b\}$ as a participant's baseline probability of being hired without investment, a Green participant's baseline probability of being hired (p_g) is 0.65; a Blue participant's baseline probability of being hired (p_b) is 0.35. In other words, a Green participant's baseline probability of scoring either an A (0.35) or a B (0.3), and a Blue participant's baseline probability of scoring an A (0.35). Recall that investing I_i increases a participant's chance of being the top performer by 1.5%, incorporating the endogenous investment decisions, a green participant's probability of being hired (p) can thus be expressed as: $(0.5 + 0.015I_i) \times 1 + (1 - 0.5 - 0.015I_i) \times 0.3 = 0.65 + 0.0105I_i$. For a Blue participant, they are only hired when they have a score of A. Their probability of being hired (p) can therefore be expressed as: $(0.5 + 0.015I_i) \times 0.7 = 0.35 + 0.0105I_i$. Subsequently, we can express participant i 's probability of being hired as a function of their baseline hiring probability p_0 and their investment (in ECU) I_i as:

$$p = P(\text{hired}) = p_0 + 0.0105I_i$$

The expected payoff (in ECU) of investing I_i in the Control treatment (μ_C) can thus be expressed as:


$$\begin{aligned}\mu_C &= p(100 - I_i) + (1 - p)(-I_i) \\ &= 100p - I_i\end{aligned}$$

which is then straightforward to calculate that the expected payoff for a Green participant is $65 + 0.05I_i$, and $35 + 0.05I_i$ for a Blue participant. This investment–qualification relationship is calibrated so that the marginal benefit of each additional ECU invested exceeds its cost. Across both groups, participants' expected payoffs are maximized when they invest the full allowable amount of 20 ECUs.

Despite that the marginal payoff of investing is the same across the two groups, the stricter hiring rule faced by the Blue participants makes their return on investment riskier. Figure A1 visualizes the probability of being hired, the expected payoffs, and the variance of payoffs in the Control treatment. As is made apparent in the figure, Blue participants face higher variance in payoff across all investment amounts compared to Green participants.

This informs our design of the “No Risk” treatment, where participants only pay their chosen investment cost if they are successfully hired, reducing the variance associated with investing in improving one's outcome. Given that participants do not incur investment costs if they're not hired in the “No Risk” treatment, their expected payoff μ_{NR} can be expressed as $p(100 - I_i) + (1 - p)(0) = p(100 - I_i)$. Let VAR_C and VAR_{NR} denote the variance of payoff in the Control and the “No Risk” treatment. We can express VAR_C as $p(100 - I_i - \mu_C)^2 + (1 - p)(-I_i - \mu_C)^2$; VAR_{NR} as $p(100 - I_i - \mu_{NR})^2 + (1 - p)(0 - \mu_{NR})^2$. Given that p is no greater than 1, μ_{NR} is strictly greater than μ_C , VAR_C is therefore also strictly higher than VAR_{NR} .

Figure A1: Hiring probability and payoff variance by group in the Control treatment



Note: Investment level in ECU is on the x-axis, probability of being hired is on the y-axis. The green (blue dash) curve and shade are the Green (Blue) participant's hiring probability and variance of payoff.

Figure A2: “Successful Role Models” and “Unsuccessful Role Models” information screens

For your reference, before you make your investment decision, here is some information about previous Blue participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	A	Yes
Participant B	Blue	20	A	Yes
Participant C	Blue	15	A	Yes
Participant D	Blue	10	A	Yes
Participant E	Blue	5	A	Yes

For your reference, before you make your investment decision, here is some information about previous Blue participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	B	No
Participant B	Blue	20	C	No
Participant C	Blue	15	B	No
Participant D	Blue	13	C	No
Participant E	Blue	10	C	No

For your reference, before you make your investment decision, here is some information about previous Green participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	A	Yes
Participant B	Green	20	B	Yes
Participant C	Green	15	A	Yes
Participant D	Green	12	A	Yes
Participant E	Green	11	A	Yes

For your reference, before you make your investment decision, here is some information about previous Green participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	C	No
Participant B	Green	11	C	No
Participant C	Green	10	C	No
Participant D	Green	10	C	No
Participant E	Green	10	C	No

Note: Top left: “Successful Role Models” screen for the Blue group; top right: “Successful Role Models” screen for the Green group; bottom left: “Unsuccessful Role Models” screen for the Blue group; bottom right: “Unsuccessful Role Models” screen for the Green group.

Figure A3: “No Risk” treatment reminders on decision page

How much ECU would you like to invest to improve your probability of becoming a top performer? Please input below any whole number from 0 to 20. Reminder: You are in the Blue Group. The hiring algorithm only hires a Blue applicant with a score of A. Your investment cost is only deducted from your earnings if you are hired. If not, you keep the full base bonus of 20 ECU.	How much ECU would you like to invest to improve your probability of becoming a top performer? Please input below any whole number from 0 to 20. Reminder: You are in the Green Group. The hiring algorithm hires a Green applicant with either a score of A or B. Your investment cost is only deducted from your earnings if you are hired. If not, you keep the full base bonus of 20 ECU.
--	---

Note: left: “No Risk” reminder for the Blue group; right: “No Risk” reminder for the Green group.

A.3 Appendix Tables

Table A1: Summary statistics by group

Variable	Blue	Green	Test differences
Age	46.24	45.34	(0.502)
Risk attitude	6.51	6.48	(0.888)
Trust in major tech companies	5.61	5.83	(0.276)
Male	48.5%	50.2%	(0.701)
Employed	69.3%	69.5%	(0.961)
Prior experience with AI or algorithm	36.1%	35.5%	(0.874)
Education: Some high school or less	1.5%	0.7%	(0.447)
Education: High school diploma or GED	14.6%	12.2%	(0.479)
Education: Some college, but no degree	16.4%	16.8%	(0.985)
Education: Associates or technical degree	12.8%	15.8%	(0.376)
Education: Bachelor’s degree	35.4%	36.2%	(0.915)
Education: Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, etc.)	19.3%	18.3%	(0.833)
Income: Less than 25k	25.5%	22.9%	(0.538)
Income: 25-50k	25.5%	25.4%	(1.000)
Income: 50-75k	16.1%	22.9%	(0.053)
Income: 75-100k	13.5%	15.8%	(0.526)
Income: 100-125k	5.5%	4.3%	(0.658)
Income: Greater than or equal to 125k	13.9%	8.6%	(0.068)
Observations	274	279	

Age is measured in years as a continuous variable. Risk attitude is recorded on a 0–10 scale, with higher scores indicating greater willingness to take risks. Trust in major tech companies is based on responses to the question: “In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way?”, also on a 0–10 scale. Male, Employed, and Prior experience with AI or algorithm are binary variables coded as 1 if participants identify as male, report being employed (full-time or part-time), or have previously submitted an application believed to be screened by an AI system or algorithm. The column labeled Test differences reports p-values from two-sample t-tests for continuous variables and difference-in-means tests for binary variables, and from chi-squared or Fisher’s exact tests for categorical variables (Education and Income).

Table A2: Summary statistics by treatment

Variable	Control	Successful RM	Unsuccessful RM	No Risk	Test differences
Age	45.62	46.03	46.14	45.20	(0.914)
Risk attitude	6.53	6.41	6.48	6.60	(0.872)
Trust in major tech companies	5.80	5.66	5.65	5.89	(0.628)
Male	46.5%	51.4%	52.5%	44%	(0.189)
Employed	68.4%	70.8%	69.9%	68.7%	(0.921)
Prior experience with AI or algorithm	36.7%	35.1%	35.7%	35.2%	(0.977)
Education: Associates or technical degree	14.2%	14.2%	14.9%	12.6%	(0.920)
Education: Bachelor's degree	35.6%	37.5%	34.2%	36.3%	(0.845)
Education: Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, etc.)	18.5%	18.2%	19.3%	19.8%	(0.968)
Education: High school diploma or GED	13.8%	12.6%	14%	13.2%	(0.958)
Education: Some college, but no degree	17.1%	16%	16.5%	17.6%	(0.967)
Education: Some high school or less	0.7%	1.5%	1.2%	0.5%	(0.731)
Income: 100-125k	3.6%	6.5%	5.9%	2.2%	(0.102)
Income: 25-50k	23.6%	27.7%	27.3%	21.4%	(0.324)
Income: 50-75k	23.3%	16.6%	16.5%	24.7%	(0.026)
Income: 75-100k	14.5%	14.2%	14.9%	15.4%	(0.984)
Income: Greater than or equal to 125k	11.3%	10.8%	10.9%	12.6%	(0.925)
Income: Less than 25k	23.6%	24.3%	24.5%	23.6%	(0.992)
Observations	275	325	322	182	

Note: Age is measured in years as a continuous variable. Risk attitude is recorded on a 0–10 scale, with higher scores indicating greater willingness to take risks. Trust in major tech companies is based on responses to the question: “In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way?”, also on a 0–10 scale. Male, Employed, and Prior experience with AI or algorithm are binary variables coded as 1 if participants identify as male, report being employed (full-time or part-time), or have previously submitted an application believed to be screened by an AI system or algorithm. The column labeled Test differences reports p-values from one-way ANOVA F-tests of equality across the four treatment groups for continuous and binary variables, and from chi-squared tests (or Fisher’s exact tests when expected cell counts are small) for categorical variables (Education and Income).

Table A3: Treatment effects by group

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group	2.317** (0.929)	2.172** (0.944)	2.014** (0.943)
Successful RM (Green)	0.503 (0.847)	0.481 (0.854)	0.286 (0.861)
Unsuccessful RM (Green)	0.152 (0.889)	0.087 (0.895)	−0.130 (0.909)
No Risk (Green)	1.216 (0.750)	1.212 (0.760)	1.207 (0.763)
Blue x Successful RM	0.551 (1.145)	0.735 (1.175)	0.834 (1.176)
Blue x Unsuccessful RM	−0.672 (1.232)	−0.405 (1.257)	−0.303 (1.261)
Blue x No Risk	2.264** (1.046)	2.206** (1.054)	2.187** (1.047)
Constant	11.744*** (0.678)	11.445*** (1.591)	12.897*** (2.284)
Round Control	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.052	0.074	0.112

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A4: Effect of group (Blue vs. Green) on investment in control treatment (excluding participants with duration time outside the 10th to 90th percentile)

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group (vs Green)	2.311** (1.038)	2.666** (1.116)	2.640** (1.145)
Constant	12.409*** (0.891)	7.327** (2.864)	7.831* (4.264)
Round Control	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	219	219	219
R ²	0.051	0.197	0.307

Note: This table reports OLS regressions estimating the effect of being assigned to the Blue group (relative to the Green group) on investment decisions in the Control treatment only. The sample excludes participants with completion times below the 10th percentile or above the 90th percentile, following the pre-registered robustness specification. The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A5: Treatment effects by group (excluding participants with duration time outside the 10th to 90th percentile)

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group	2.435** (1.039)	2.251** (1.053)	2.121** (1.037)
Successful RM (Green)	0.971 (0.948)	0.914 (0.950)	0.799 (0.954)
Unsuccessful RM (Green)	0.617 (0.970)	0.493 (0.973)	0.388 (0.983)
No Risk (Green)	1.727** (0.816)	1.756** (0.821)	1.643** (0.824)
Blue x Successful RM	0.336 (1.286)	0.374 (1.307)	0.560 (1.305)
Blue x Unsuccessful RM	−1.449 (1.381)	−1.336 (1.399)	−1.149 (1.387)
Blue x No Risk	2.137* (1.169)	2.025* (1.180)	2.192* (1.173)
Constant	11.681*** (0.742)	10.172*** (1.790)	11.454*** (2.440)
Round Control	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	886	886	886
R ²	0.057	0.093	0.142

Note: This table reports OLS regressions estimating treatment effects by group, excluding participants with completion times below the 10th percentile or above the 90th percentile, following the pre-registered robustness specification. The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A6: Effect of group (Blue vs. Green) on investment in control treatment using only first-round data

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group (vs Green)	1.939 (1.259)	1.967 (1.384)	2.167 (1.413)
Constant	12.554*** (0.921)	8.558*** (3.063)	5.265 (4.805)
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	140	140	140
R ²	0.017	0.154	0.358

Note: This table reports OLS regressions estimating the effect of group assignment (Blue vs. Green) on investment in the Control treatment, using only first-round observations as specified in the pre-registration. The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A7: Treatment effects using only first-round data

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue Group	1.939 (1.209)	1.458 (1.245)	1.422 (1.243)
Successful RM (Green)	1.087 (1.256)	0.739 (1.289)	0.568 (1.297)
Unsuccessful RM (Green)	−1.878 (1.213)	−2.079* (1.248)	−2.439* (1.264)
No Risk (Green)	−0.475 (1.205)	−0.653 (1.243)	−0.895 (1.271)
Blue x Successful RM	0.302 (1.733)	1.208 (1.800)	0.987 (1.806)
Blue x Unsuccessful RM	0.357 (1.693)	0.825 (1.757)	1.213 (1.765)
Blue x No Risk	3.048* (1.726)	3.305* (1.789)	3.261* (1.810)
Constant	12.554*** (0.885)	12.491*** (1.895)	12.562*** (2.591)
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	553	553	553
R ²	0.069	0.099	0.154

Note: This table reports OLS regressions estimating treatment effects by group using only first-round observations, following the pre-registered robustness specification. The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). Column (2) adds demographic controls: age (continuous), gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Robust standard errors, clustered at the participant level, are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A8: Treatment effects by age

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.654*** (0.547)	2.651*** (0.555)	2.539*** (0.545)
Successful RM	0.061** (0.031)	0.056* (0.033)	0.062* (0.033)
Unsuccessful RM	3.568** (1.796)	3.418* (1.815)	3.272* (1.815)
No Risk	1.860 (1.842)	1.855 (1.856)	1.681 (1.837)
Age	3.251** (1.645)	3.260* (1.670)	2.995* (1.665)
Age $\times$ Successful RM	−0.061 (0.037)	−0.056 (0.038)	−0.056 (0.038)
Age $\times$ Unsuccessful RM	−0.045 (0.038)	−0.043 (0.038)	−0.043 (0.038)
Age $\times$ No Risk	−0.020 (0.034)	−0.021 (0.035)	−0.016 (0.035)
Constant	8.771*** (1.543)	9.798*** (1.905)	11.275*** (2.498)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.054	0.073	0.111

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: gender (binary), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Age is measured in years as a continuous variable. Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A9: Treatment effects by gender

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.657*** (0.548)	2.628*** (0.555)	2.528*** (0.546)
Successful RM	-1.349 (0.934)	-1.526 (0.952)	-1.693* (0.948)
Unsuccessful RM	-0.117 (0.777)	-0.023 (0.779)	-0.074 (0.782)
No Risk	-1.233 (0.865)	-1.124 (0.867)	-1.189 (0.865)
Gender	1.551** (0.682)	1.568** (0.691)	1.543** (0.686)
Gender $\times$ Successful RM	1.859 (1.148)	1.842 (1.149)	1.629 (1.153)
Gender $\times$ Unsuccessful RM	2.133* (1.235)	2.082* (1.235)	1.861 (1.232)
Gender $\times$ No Risk	1.626 (1.058)	1.520 (1.079)	1.542 (1.087)
Constant	12.205*** (0.676)	12.041*** (1.491)	13.883*** (2.114)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.052	0.074	0.111

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), education level (ordered factor), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Gender is coded as 1 if the participant identifies as male, 0 if the participant identifies as female or non binary/third-gender. Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A10: Treatment effects by education

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.778*** (0.548)	2.716*** (0.552)	2.629*** (0.544)
Successful RM	2.382 (3.785)	1.772 (3.955)	1.661 (3.827)
Unsuccessful RM	−2.406 (3.490)	−2.877 (3.613)	−2.953 (3.473)
No Risk	−0.365 (2.515)	−0.330 (2.592)	−0.248 (2.497)
Education	1.012 (1.599)	0.676 (1.632)	0.581 (1.575)
Education × Successful RM	1.899 (1.210)	1.873 (1.190)	1.954* (1.168)
Education × Unsuccessful RM	1.000 (1.124)	1.132 (1.148)	0.993 (1.098)
Education × No Risk	0.693 (1.310)	0.826 (1.335)	0.641 (1.264)
Constant	11.139*** (1.147)	10.638*** (1.754)	11.972*** (2.403)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.071	0.087	0.126

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), employment status (binary), and income bracket (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Education is an ordinal variable ranging from 1 (some high school or less) to 6 (graduate or professional degree). Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A11: Treatment effects by income

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.651*** (0.549)	2.646*** (0.546)	2.612*** (0.545)
Successful RM	0.766 (0.581)	0.751 (0.582)	0.756 (0.581)
Unsuccessful RM	-0.197 (0.618)	-0.223 (0.620)	-0.228 (0.619)
No Risk	2.288*** (0.522)	2.294*** (0.520)	2.271*** (0.519)
Income	0.231 (0.212)	0.225 (0.213)	0.218 (0.210)
Income $\times$ Successful RM	-0.465* (0.254)	-0.445* (0.252)	-0.447* (0.250)
Income $\times$ Unsuccessful RM	-0.328 (0.274)	-0.314 (0.273)	-0.307 (0.273)
Income $\times$ No Risk	0.028 (0.237)	0.030 (0.238)	0.037 (0.237)
Constant	14.322*** (0.941)	13.319*** (1.634)	11.749*** (1.978)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.078	0.092	0.134

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education (ordered factor), and employment status (binary). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Income is a participant’s gross annual salary, an ordinal variable ranging from 1 (less than 25,000 USD) to 13 (greater than/equal to 300,000 USD). Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A12: Treatment effects by employment status

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.648*** (0.549)	2.564*** (0.558)	2.453*** (0.547)
Successful RM	−0.169 (1.041)	−0.469 (1.092)	−0.429 (1.080)
Unsuccessful RM	0.745 (1.107)	0.597 (1.110)	0.438 (1.109)
No Risk	−0.029 (1.144)	−0.161 (1.147)	−0.283 (1.147)
Employment	3.721*** (0.933)	3.650*** (0.925)	3.740*** (0.921)
Employment × Successful RM	0.054 (1.284)	0.368 (1.295)	0.388 (1.300)
Employment × Unsuccessful RM	−0.224 (1.356)	0.067 (1.365)	0.013 (1.366)
Employment × No Risk	−2.070* (1.137)	−1.997* (1.133)	−2.163* (1.127)
Constant	11.705*** (0.916)	12.271*** (1.741)	13.406*** (2.459)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.052	0.068	0.107

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education (ordered factor), and income (ordered factor). Column (3) further includes individual-level covariates: risk attitude (self-reported on a 0–10 scale), prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Employment status is coded as 1 if the participant reports to be working full-time or part-time, 0 if the participant reports to be unemployed and looking for work, a homemaker or stay-at-home parent, a student or retired. Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A13: Treatment effects by risk attitude

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.587*** (0.546)	2.516*** (0.554)	2.439*** (0.552)
Successful RM	-4.240 (3.358)	-4.496 (3.155)	-4.535 (3.346)
Unsuccessful RM	-3.616 (2.658)	-3.864 (2.569)	-3.116 (2.683)
No Risk	0.385 (2.299)	0.539 (2.257)	1.052 (2.385)
Risk attitude	-1.720 (2.475)	-1.540 (2.391)	-1.293 (2.526)
Risk attitude × Successful RM	0.311 (2.220)	0.544 (2.165)	1.080 (2.361)
Risk attitude × Unsuccessful RM	-2.319 (2.021)	-2.374 (1.967)	-1.624 (2.140)
Risk attitude × No Risk	-1.835 (2.103)	-2.195 (2.078)	-1.673 (2.247)
Constant	13.008*** (1.819)	12.246*** (2.342)	12.486*** (2.776)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.081	0.107	0.126

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education (ordered factor), employment status (binary), and income (ordered factor). Column (3) further includes individual-level covariates: prior experience with algorithmic screening (binary), and trust in major technology companies (0–10 scale). Risk attitude is the participant’s response on a 0 to 10 scale to the question: “In general, how willing are you to take risks? (0 = Completely unwilling to take risks; 10 = Very willing to take risks).” Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A14: Treatment effects by trust in major tech companies

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.674*** (0.558)	2.610*** (0.560)	2.546*** (0.551)
Successful RM	0.182 (2.779)	−0.586 (2.824)	−0.933 (2.805)
Unsuccessful RM	0.508 (2.548)	−0.214 (2.554)	−0.468 (2.523)
No Risk	0.747 (2.319)	−0.260 (2.285)	−0.644 (2.307)
Trust	0.530 (2.079)	−0.070 (2.057)	−0.166 (2.080)
Trust × Successful RM	−2.996 (2.199)	−3.787* (2.140)	−4.026* (2.180)
Trust × Unsuccessful RM	−0.920 (2.314)	−1.727 (2.280)	−2.401 (2.337)
Trust × No Risk	−2.365 (2.201)	−3.115 (2.153)	−3.711* (2.193)
Constant	12.363*** (1.873)	12.573*** (2.420)	13.184*** (2.738)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Prior Experience	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.074	0.100	0.124

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education (ordered factor), employment status (binary), and income (ordered factor). Column (3) further includes individual-level covariates: risk attitude (0-10 scale), and prior experience with algorithmic screening (binary). Trust is the participant's response on a 0 to 10 scale to the question: "In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way? (0 = No trust at all; 10 = Complete trust)." Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A15: Treatment effects by prior experience with AI/algorithms

	<i>Dependent variable:</i>		
	Investment		
	(1)	(2)	(3)
Blue	2.660*** (0.549)	2.631*** (0.554)	2.573*** (0.547)
Successful RM	−0.866 (0.959)	−0.922 (0.979)	−1.276 (0.996)
Unsuccessful RM	0.624 (0.720)	0.702 (0.729)	0.522 (0.716)
No Risk	−0.102 (0.788)	−0.050 (0.799)	−0.249 (0.780)
Prior experience	2.255*** (0.643)	2.228*** (0.653)	2.206*** (0.657)
Prior experience × Successful RM	0.374 (1.195)	0.402 (1.202)	0.685 (1.209)
Prior experience × Unsuccessful RM	−0.273 (1.260)	−0.182 (1.265)	0.080 (1.263)
Prior experience × No Risk	0.118 (1.108)	0.112 (1.114)	0.095 (1.114)
Constant	11.875*** (0.651)	11.928*** (1.635)	13.819*** (2.245)
Round Fixed Effects	Yes	Yes	Yes
Demographic Controls	No	Yes	Yes
Risk Attitude	No	No	Yes
Trust in Tech	No	No	Yes
Observations	1,106	1,106	1,106
R ²	0.051	0.073	0.106

Note: The dependent variable is the number of investment units allocated (ranging from 0 to 20 inclusive). All models include round fixed effects. Column (2) adds demographic controls: age (continuous), gender (binary), education (ordered factor), employment status (binary), and income (ordered factor). Column (3) further includes individual-level covariates: risk attitude (0-10 scale), and trust in major technology companies (0–10 scale). Prior experience is coded as 1 if the participant responds to the question: “Have you ever applied for a job or other opportunity where you believe your application was screened or reviewed by an AI system or algorithm? (For example: resume filtering by applicant tracking systems used by companies and recruitment agencies; or credit checks by systems like FICO Score.)” by selecting the option: “Yes, I have submitted an application where I believe my application was screened by an AI system or algorithm.” 0 if otherwise. Robust standard errors clustered at the participant level in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table A16: Hiring rates by group and treatment

Treatment	Blue	Green	Gap (Green - Blue)
Control	77.3	50.7	26.6
Successful RM	83.9	47.0	36.9
Unsuccessful RM	80.1	47.8	32.3
No Risk	78.9	48.3	30.7

Note: Entries report the percentage of participants hired by the algorithm. The gap is calculated as the hiring rate of Green participants minus the hiring rate of Blue participants. Hiring decisions follow the group-specific hiring rule described in Table 3.

B Full Survey



Welcome & Consent

Welcome and Study Information

You are invited to take part in a short research study on decision-making, conducted by a researcher from the School of Economics at University College Dublin.

What you'll do:

- You will be given some information about a scenario and then you will be asked to make a number of decisions
- You have the chance to earn an additional bonus depending on your decisions
- Answer a brief demographic questionnaire

Participation details:

- You must be at least **18 years old**
- The study takes about **10 minutes**
- On top of the participation fee, you will earn a bonus between **\$0 USD** to **\$1.2 USD** based on your decisions

Data & Privacy:

- Your responses are **anonymous** and **confidential**
- No personal identifiers (e.g. name, email, IP) will be collected
- Data will be stored securely and handled according to GDPR and university policies

Voluntary Participation:

- Participation is voluntary — you can withdraw at any time before submitting
- After submission, data is anonymized and cannot be withdrawn

Questions?

You can contact the researcher at yung-shiang.yang@ucdconnect.ie.

By clicking the button below, you confirm that:

- You are at least 18 years old
- You understand your participation is voluntary
- You may withdraw at any time before submission

I consent, begin the study

I do not consent, I do not wish to participate

Prolific ID

What is your Prolific ID? Please note that this response should auto-fill with the correct ID

`#{e://Field/PROLIFIC_PID}`

Payment Information

Payment Information

About ECU:

You'll earn and spend points called **ECU (Experimental Currency Units)**, **100 ECU = \$1.0 USD**, this is used to calculate your bonus.

You will start with:

20 ECU + a chance to earn an additional **100 ECU** based on your decision and outcome in a randomly selected round in the experiment.

Make your decisions carefully.

Scenario Instructions - Control

Scenario Overview

- You are applying for a job where a **hiring algorithm** decides who gets hired.

- There are two groups of applicants: **Green** and **Blue**. You've been assigned to the **`#{e://Field/GroupColor}`**.
-

Goal of the Algorithm

- Each applicant will either become a **top** or an **average** performer in the job.
 - The algorithm only wants to hire applicants who will become **top performers** in the job, but it can't tell who they are with certainty.
 - Therefore, it asks that every applicant take a **pre-employment test** — a signal, though not a perfect one, of future performance.
-

Algorithm's Belief

- The algorithm uses both your **test score** and your **group identity** to decide whether to hire you.
 - It believes **Green** applicants are more likely to become top performers than **Blue** applicants — even though in reality, **everyone starts with a 50%** chance of becoming a top performer.
-

Investment Option

- You can invest any amount from your base bonus of **20 ECU** to boost your chance of becoming a **top performer**.
 - **1 ECU = +1.5%** boost. So, **10 ECU = +15%**; **20 ECU = +30%**.
 - Your total chance of becoming a top performer = **50% + investment boost**.
-

The Pre-Employment Test

- If you're **a top performer**, you have 70% chance of scoring **A** (the highest score), 30% chance of scoring **B**, and 0% chance of scoring **C** (the lowest score).
 - If you're **an average performer**, you have 0% chance of scoring **A**, 30% chance of scoring **B**, and 70% chance of scoring **C**.
-

The Algorithm's Hiring Rule

- **Green:** Hired if your score is **A or B**
 - **Blue:** Hired only if your score is **A**
-

Bonus

- **Base bonus:** 20 ECU
 - **If hired:** +100 ECU
 - **Important: Your investment cost is deducted from the base bonus whether you're hired or not.**
-

Examples

- Invest 10 ECU & hired → Bonus = $20 - 10 + 100 = \mathbf{110\ ECU}$
 - Invest 10 ECU & not hired → Bonus = $20 - 10 + 0 = \mathbf{10\ ECU}$
 - Invest 0 ECU & hired → Bonus = $20 - 0 + 100 = \mathbf{120\ ECU}$
-

 **You'll complete 1 practice and 2 formal rounds. Your bonus depends on 1 randomly chosen formal round — so decide carefully!**

Scenario Instructions - No Risk

Scenario Overview

- You are applying for a job where a **hiring algorithm** decides who gets hired.
 - There are two groups of applicants: **Green** and **Blue**. You've been assigned to the **#{e://Field/GroupColor}**.
-

Goal of the Algorithm

- Each applicant will either become a **top** or an **average performer** in the job.
 - The algorithm only wants to hire applicants who will become **top performers** in the job, but it can't tell who they are with certainty.
 - Therefore, it asks that every applicant take a **pre-employment test** — a signal, though not a perfect one, of future performance.
-

Algorithm's Belief

- The algorithm uses both your **test score** and your **group identity** to decide whether to hire you.
 - It believes **Green** applicants are more likely to become top performers than **Blue** applicants — even though in reality, **everyone starts with a 50%** chance of becoming a top performer.
-

Investment Option

- You can invest any amount from your base bonus of **20 ECU** to improve your chance of becoming a top performer.
- **1 ECU = +1.5%** chance. So, **10 ECU = +15%**; **20 ECU = +30%**.

- Your total chance of becoming a top performer = **50% + investment boost.**
-

The Pre-Employment Test

- If you're a top performer, you have 70% chance of scoring **A** (the highest score), 30% chance of scoring **B**, and 0% chance of scoring **C** (the lowest score).
 - If you're an average performer, you have 0% chance of scoring **A**, 30% chance of scoring **B**, and 70% chance of scoring **C**.
-

Hiring Rule

- **Green:** Hired if your score is **A or B**
 - **Blue:** Hired only if your score is **A**
-

Bonus

- **Base bonus:** 20 ECU
 - **If hired:** +100 ECU
 - **Important: Your investment cost is only deducted from the base bonus if you're hired.**
-

Examples

- Invest 10 ECU & hired → Bonus = $20 - 10 + 100 =$ **110 ECU**
 - Invest 10 ECU & not hired → Bonus = $20 - 0 + 0 =$ **10 ECU**
 - Invest 0 ECU & hired → Bonus = $20 - 0 + 100 =$ **120 ECU**
-

 **You'll complete 1 practice and 2 formal rounds. Your bonus depends on 1 randomly chosen formal round — so decide carefully!**

Attention Check

Attention Check

To ensure you've understood the instructions, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

Please drag and drop the following events into the correct order based on how the study works:

Test score is simulated

You are randomly assigned to a group

You decide how much to invest

The hiring algorithm makes a hiring decision

What is your chance of becoming a top performer in the job if you invest the maximum amount of 20 ECU?

(Hint: Everyone starts with a 50% chance of becoming a top performer. Each 1 ECU you invest increases your chance by 1.5 percentage points.)

65%

70%

75%

80%

Suppose, following the investment, you end up becoming a top performer. What are your chances of scoring A, B, C, respectively?

70%, 30%, 0%

0%, 30%, 70%

100%, 0%, 0%

Suppose you're a **Green applicant with a score of B. Will you be hired?**

*(Hint: The algorithm hires a **Green** applicant with either an **A** or **B**, only hires a **Blue** applicant with an **A**.)*

Yes

No

Suppose you're hired following your investment of 20 ECU. What will be the total amount of your bonus?

(Hint: You start off with a base bonus of 20 ECU, which you can invest at your discretion. You also get an additional bonus of 100 ECU if you're hired.)

120 ECU

100 ECU

0 ECU

20 ECU

Blue SRM

For your reference, before you make your investment decision, here is some information about previous **Blue participants who have completed this task:**

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	A	Yes
Participant B	Blue	20	A	Yes
Participant C	Blue	15	A	Yes
Participant D	Blue	10	A	Yes
Participant E	Blue	5	A	Yes



Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Green SRM

For your reference, before you make your investment decision, here is some information about previous **Green participants who have completed this task:**

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	A	Yes
Participant B	Green	20	B	Yes
Participant C	Green	15	A	Yes

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant D	Green	12	A	Yes
Participant E	Green	11	A	Yes

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Blue URM

For your reference, before you make your investment decision, here is some information about previous **Blue participants who have completed this task:**

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	B	No
Participant B	Blue	20	C	No
Participant C	Blue	15	B	No
Participant D	Blue	13	C	No
Participant E	Blue	10	C	No

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the

algorithm.

hired

not hired

Green URM

For your reference, before you make your investment decision, here is some information about previous **Green** participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	C	No
Participant B	Green	11	C	No
Participant C	Green	10	C	No
Participant D	Green	10	C	No
Participant E	Green	10	C	No

 Attention Check

To ensure you’ve read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Practice Round



Practice Round

This round is just for practice and will not affect your final bonus. It is designed to help you understand how the investment and test process work.

How much ECU would you like to invest to improve your probability of becoming a top performer? Please input below any whole number from 0 to 20.

Reminder: You are in the \${e://Field/GroupColor}.

\${e://Field/GroupReminder}

Your investment cost is deducted from your earnings whether you're hired or not.

Practice Simulation

Thanks for making your decision!

Your test result is now simulating...



Click  to reveal your results.

Practice Results



Practice Round Result

Here's how your decision played out in the practice round (this will not affect your final bonus):



You invested: $\${e://Field/PracticeInvestment}$



Your probability of becoming a top performer after investment is: $\${e://Field/PracticeProb}\%$



Your test score: $\${e://Field/PracticeTestScoreLabel}$



Hiring decision: You are $\${e://Field/PracticeHired}$ by the hiring algorithm.



Your (hypothetical) bonus in the practice round: 20 -

$\${e://Field/PracticeInvestmentDisplay} + \${e://Field/PracticeWage} = \${e://Field/PracticeEarnings}$ ECU

Note: This was just a practice round. You'll now complete **2 formal rounds**, and your bonus will be based on *one randomly selected formal round*.

Formal Round

✓ Formal Round 1 of 2

This is the first of two formal rounds. If this is the randomly chosen round, your decision here WILL affect your final bonus.

How much ECU would you like to invest to improve your probability of becoming a top performer? Please input below any whole number from 0 to 20.

Reminder: You are in the $\{e://Field/GroupColor\}$.
 $\{e://Field/GroupReminder\}$

Simulation

Thanks for making your decision!

Your test result is now simulating...



Click  to reveal your results.

Results

Result of Round 1 of 2

This was the first of two formal rounds. If this is the randomly chosen round, the outcome below WILL determine your final bonus.

 **You invested:** $\${e://Field/Investment}$

 **Your probability of becoming a top performer after investment is:** $\${e://Field/Prob}\%$

 **Your test score:** $\${e://Field/TestScoreLabel}$

 **Hiring decision:** You are $\${e://Field/Hired}$ by the hiring algorithm.

 **Your bonus:** $20 - \${e://Field/InvestmentDisplay} + \${e://Field/Wage} = \${e://Field/Earnings}$ ECU

You've completed the first formal round. You'll now proceed to the second and final formal round.

Green LRM 2

For your reference, before you make your investment decision in the final round, here is some information about previous **Green participants who have completed this task:**

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	A	Yes
Participant B	Green	20	B	Yes
Participant C	Green	15	A	Yes
Participant D	Green	12	A	Yes
Participant E	Green	11	A	Yes

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Blue LRM 2

For your reference, before you make your investment decision in the final round, here is some information about previous **Blue** participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	A	Yes
Participant B	Blue	20	A	Yes
Participant C	Blue	15	A	Yes
Participant D	Blue	10	A	Yes
Participant E	Blue	5	A	Yes

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Green URM 2

For your reference, before you make your investment decision in the final round, here is some information about previous **Green** participants who have completed this task:

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Green	20	C	No
Participant B	Green	11	C	No
Participant C	Green	10	C	No
Participant D	Green	10	C	No
Participant E	Green	10	C	No

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Blue URM 2

For your reference, before you make your investment decision in the final round, here is some information about previous **Blue participants who have completed this task:**

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant A	Blue	20	B	No
Participant B	Blue	20	C	No
Participant C	Blue	15	B	No

Participant	Group	Investment (ECU)	Test Score	Hired?
Participant D	Blue	13	C	No
Participant E	Blue	10	C	No

Attention Check

To ensure you've read the information, please answer the following questions. (You will not be able to proceed unless you answer correctly.)

The previous participants shown above invested ____ ECU to improve their chances of becoming top performers.

some

zero

The previous participants shown above were ____ by the algorithm.

hired

not hired

Second Round No Risk Instructions

Attention: Change in Rules

In the following round, your investment cost is **only deducted if you're hired**. Your investment cost is **not deducted if you're not hired**. The hiring rule from the hiring algorithm remains the same.

Second Round Control Instructions

Attention: Change in Rules

In the following round, your investment cost **WILL be deducted whether you're hired or not**. The hiring rule from the hiring algorithm remains the same.

Second Formal Round

Formal Round 2 of 2

This is the second and final formal round. If this is the randomly chosen round, your decision here WILL affect your final bonus.

How much ECU would you like to invest to improve your probability of becoming a top performer? Please input below any whole number from 0 to 20.

Reminder: You are in the $\{e://Field/GroupColor\}$.
 $\{e://Field/GroupReminder\}$

Second Simulation

Thanks for making your decision!

Your test result is now simulating...



Click  to reveal your results.

Second Results

Result of Round 2 of 2

This was the second and final formal round. If this is the randomly chosen round, the outcome below WILL determine your final bonus.

 **You invested:** $\{e://Field/Investment2\}$

 **Your probability of becoming a top performer after investment is:** $\{e://Field/Prob2\}\%$

 **Your test score:** $\{e://Field/TestScoreLabel2\}$

 **Hiring decision:** You are $\{e://Field/Hired2\}$ by the hiring algorithm.

 **Your bonus:** $20 - \{e://Field/InvestmentDisplay2\} + \{e://Field/Wage2\} = \{e://Field/Earnings2\}$ ECU

This concludes the decision-making part of the study. You'll now answer a short demographic questionnaire.

Demographics

Please answer the following short questions.

Your responses are anonymous and will be used for research purposes only. You will not be personally identifiable from your answers.

What is your age?

(Please enter in years)

What is your gender?

Male

Female

Non-binary / third gender

Prefer not to say

What is the highest level of education you have completed?

Some high school or less

High school diploma or GED

Some college, but no degree

Associates or technical degree

Bachelor's degree

Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, etc.)

Prefer not to say

What best describes your employment status?

Working full-time

Working part-time

Unemployed and looking for work

A homemaker or stay-at-home parent

Student

Retired

Other (please specify in the textbox below)

Please estimate your gross annual salary.

Less than 25,000 USD

Greater than/equal to 25,000 USD and less than 50,000 USD

Greater than/equal to 50,000 USD and less than 75,000 USD

Greater than/equal to 75,000 USD and less than 100,000 USD

Greater than/equal to 100,000 USD and less than 125,000 USD

Greater than/equal to 125,000 USD and less than 150,000 USD

Greater than/equal to 150,000 USD and less than 175,000 USD

Greater than/equal to 175,000 USD and less than 200,000 USD

Greater than/equal to 200,000 USD and less than 225,000 USD

Greater than/equal to 225,000 USD and less than 250,000 USD

Greater than/equal to 250,000 USD and less than 275,000 USD

Greater than/equal to 275,000 USD and less than 300,000 USD

Greater than/equal to 300,000 USD

Prefer not to say

In general, how willing are you to take risks?

0 = Completely unwilling to take risks 10 = Very willing to take risks

0 ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐

Have you ever applied for a job or other opportunity where you believe your application was screened or reviewed by an AI system or algorithm?

(For example: resume filtering by applicant tracking systems used by companies and recruitment agencies; or credit checks by systems like FICO Score.)

Yes, I have submitted an application where I believe my application was screened by an AI system or algorithm

No, I have submitted applications but do not believe an AI system or an algorithm was ever used to screen them

I have never applied for a job or an opportunity that would involve screening

Not sure

Prefer not to say

If so, in which contexts have you interacted with the AI system or algorithm?

Select all that apply:

Job applications or hiring platforms

Credit scoring or loan applications

Other (please specify in the textbox below)

In general, how much do you trust major technology companies to make decisions that affect you in a fair and unbiased way?

0 = No trust at all 10 = Complete trust

0 ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐

In the experiment, what factor(s) below motivated your investment choice? Select all that apply.

My group identity (whether I was assigned to the **Green** or **Blue** group)

Whether I have to pay for the investment cost when I'm not hired

How much previous participants invested

Whether previous participants were hired

The starting amount I had available to invest (20 ECU)

The potential bonus I could get if hired (100 ECU)

Other (please specify in the textbox below)

Please let us know if you experienced any issue in the study, for example if anything was unclear or didn't work (optional).

End of Survey Message

Thank You for Participating!

This study investigates how people make decisions about investing in their skills when evaluated by a hiring algorithm. The hiring process in this experiment was simulated using fictional groups (**Green** and **Blue**) to study how perceived bias might affect decision-making. These group assignments were entirely hypothetical and not tied to any real-world identities.

Your bonus payment was calculated using Experimental Currency Units (ECU). The conversion rate is: **100 ECU = \$1 USD**. The final bonus amount (if any) will be transferred to your Prolific account in the next few weeks.

If you have any questions about the study or would like to learn more about the research, feel free to contact the researcher at: yung-shiang.yang@ucdconnect.ie

Please click the button below to be redirected back to Prolific and register your submission.

Alternatively, you can enter the following completion code to register your submission: **CS8K8PPI**

Powered by Qualtrics