Better Applications, Worse Matching: Artificial Intelligence and Talent Allocation*

Thomas Jungbauer†

April 8, 2026

Abstract. Generative Artificial Intelligence (AI) raises productivity inside jobs while reducing the informativeness of the screening outputs used to allocate workers across firms. In a stylized matching model, the market sorts efficiently on visible rank, but rank is an imperfect proxy for worker type. Aggregate output, decomposed into average worker output and a misallocation penalty, can therefore fall even when every worker-firm pair becomes more productive. The analysis predicts weaker screening validity, greater use of harder-to-fake assessments, socially excessive applicant AI use as well as greater reliance on probationary hiring and rematching. Empirical design suggestions to test these predictions are provided. (JEL D82, J24, J41, O33.)

Keywords: AI, Labor-Market Screening, Matching, Sorting, Talent Allocation.

*The author is grateful for helpful comments from Stefania Albanesi, David Andolfatto, Yi Chen, Benjamin Friedrich, Jie Gong, Christian Haefke, Mitchell Hoffman, Justin Johnson, Chenyang Li, Jin Li, Jeanine Miklos-Thal, Nicola Persico, Marcel Preuss, Michael Raith, Christian Schmid, Daniela Scur, Gerhard Speckbacher, Michael Waldman and Cole Williams. All remaining errors are solely the author's.
†Assistant Professor of Strategy & Business Economics, Cornell University: jungbauer@cornell.edu.

1 Introduction

Generative Artificial Intelligence (GenAI) is changing the informational foundation of hiring. Workers can now use AI to rewrite resumes, tailor cover letters, polish writing samples, generate portfolio-like artifacts, and script interview responses at negligible marginal cost. Recent reporting describes an emerging arms race in which applicants use AI to produce polished materials at scale while employers respond with additional verification, more structured assessment, and, in some settings, a return to in-person interviews (Ellis, 2024; Smith, 2025; The Economist, 2026). This development creates a theoretical possibility investigated in this paper: a technology that increases productivity throughout can lead to worse matching and potentially lower aggregate output.

A growing literature shows that GenAI can raise a worker's output when performing a task (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Dell'Acqua, McFowland III, Mollick, Lifshitz, Kellogg, Rajendran, Krayer, Candelon, and Lakhani, 2026). Hiring, however, takes place before those within-job gains are realized. Labor markets allocate workers using screening objects—resumes, cover letters, tests, portfolios, short work samples, and similar artifacts. If applicant-facing AI makes those objects more polished but less diagnostic of fit, a technology that improves performance inside realized matches need not improve who is matched to which job. In fact, recent empirical work already suggests that applicant-facing AI can weaken the signaling value of written applications and reduce screening accuracy (Galdin and Silbert, 2025; Cui, Dias, and Ye, 2025; Cowgill, Hernandez-Lagos, and Wright, 2026).

This paper develops a stylized one-to-one matching model with transferable utility that separates AI's effect on productivity inside realized matches from AI's effect on the information used to form those matches. Workers are indexed by a one-dimensional latent fit type and firms by a one-dimensional task type. The setup captures hiring environments in which pre-match screening objects generate a broadly comparable public ordering across applicants—platform markets, standardized assessments, common credential screens, and hiring funnels built around shared application artifacts. Production combines a worker-specific component that is valuable at every firm with a fit component that depends on how closely the worker's type aligns with the firm's ideal task. Before matching takes place, workers generate observable screening scores through a regime-specific screening technology. The market does not observe fit directly. It observes only each

worker's position in the cross-sectional distribution of screening outputs, represented by market rank. Screening is therefore imperfect both before and after the arrival of AI. The central question is how AI changes the relation between latent type and visible rank.

The first result characterizes assignment. Both the planner and the competitive market sort workers and firms assortatively on visible rank. The market is therefore efficient conditional on what it sees. The distortion comes from a different source, namely that rank is only an imperfect proxy for the latent dimension that governs fit. The labor-market effect of AI is not a generic breakdown of sorting. It is a change in the informational quality of the ranking on which sorting is based.

The second result gives a decomposition of aggregate output. Output equals average worker-specific productivity minus an allocative loss created by imperfect screening. That loss is pinned down by two objects only. First, the importance of exact fit in production and, second, the extent to which rank fails to track worker type. Poorer screening is therefore more costly when mismatch is more expensive, and less costly when rank remains tightly aligned with worker type. The allocative role of first-stage screening thus collapses to a single statistic measuring the informational quality of market rank.

This decomposition drives the regime comparison. Relative to the pre-AI regime, AI raises direct productivity for every worker type and reduces mismatch cost. Hence every worker–firm pair becomes more productive. However, AI may weaken the screening technology by making rank a poorer proxy for type. Under these conditions, pairwise productivity gains do not imply better allocation. Productivity works inside matches, whereas screening determines which matches are formed. When AI improves the first margin but degrades the second, the informational loss can dominate. A sufficient condition under which aggregate output falls despite higher pairwise productivity is characterized.

A complementary result distinguishes informational change from cosmetic improvement. If AI changes the appearance, fluency, or format of screening outputs without changing the ranking they induce, then assignment, output, and competitive payoffs are unchanged up to the standard normalization of transfers. The economically relevant question is therefore not whether application materials look better on average, but whether the induced visible ordering contains more or less information about worker fit.

The baseline allocation is decentralized as an assignment game. Competitive equilibria exist, but, as is common, transfers are not unique. The analysis therefore builds on the firm-

optimal equilibrium, assigning the lowest ranked worker zero rent. This normalization preserves the efficient visible-rank assignment and pins down a unique continuation value for firms, thereby providing a natural benchmark for the decentralized extensions.[1]

In real-world markets, workers choose whether and how intensively to use AI when applying for jobs. An extension of the base model thus endogenizes applicant-side AI use. Workers choose how intensively to use application AI, taking the wage schedule induced by the market as given. Application AI is privately attractive because it can improve expected market rank and thus expected pay. Yet when all workers intensify use simultaneously, the principal general-equilibrium effects are weaker screening and the direct resource cost of adoption. The decentralized equilibrium therefore resembles a prisoner's dilemma featuring excessive use of application AI relative to the social optimum. The mechanism is positional rather than productive: privately profitable efforts to climb the visible ranking erode the social value of the ranking itself.

Finally, the paper studies how weaker first-stage screening shifts selection into the employment relationship. In a two-period environment, employers learn more about incumbents after hiring. They may then separate and enter a secondary market, or they may pay ex ante for a probationary contract that preserves the option to terminate after a short evaluation period. The key object is the value of correcting matches that look promising based on first-stage information but are revealed to be weak once richer information arrives. Noisier first-stage screening raises that option value, making both rematching and probationary hiring more attractive.

The contribution is distinct from adjacent literatures. First, the model is not a signaling model in which costly actions separate worker types. The central distortion arises also under symmetric use of technology, because AI changes the mapping from public screening artifacts to fit. Second, the model is not primarily an employer-learning model either. The core inefficiency is already present before employment begins as weaker first-stage information changes who is hired. Third, the model departs from standard assignment formulations that take the public ranking as primitive. The screening technology that generates the ranking is itself endogenous or regime-dependent. The market remains efficient conditional on observables, but the observables themselves become less informative.

[1] Alternative competitive payoff selections support the same visible-rank assignment and qualitative results. The representative normalization matters quantitatively because the later sections on firm policy require a definite continuation value of a worker-firm relationship for the worker's employer.

The framework is deliberately targeted to hiring environments in which applicants are first compared through broadly common screening objects. Resumes, cover letters, standardized tests, portfolios, asynchronous interviews, and short work samples all fit this structure when they feed a common pre-match ordering. The formal analysis therefore studies how the informativeness of that ordering changes with AI. In the baseline regime comparison, the deterioration in screening is an exogenous shift in the score technology. In the endogenous-application section, it is generated by applicant choices.

In fact, the distinction between productivity and screening is especially relevant in modern hiring because most organizations evaluate workers through a sequence of filters. Low-cost artifacts determine who reaches richer and more expensive forms of assessment. A deterioration in the informational content of those artifacts therefore affects the entire hiring funnel rather than only the first screen. The same logic extends to freelance platforms, professional-service procurement, startup recruiting, and internal promotion contests, all of which rank individuals before more informative interaction occurs.

The analysis also yields four direct empirical predictions. First, applicant-facing AI makes screening less predictive of realized fit, even when applications look better on average. Second, firms should substitute toward harder-to-fake assessments and verification. Third, applicant-side AI use should be privately attractive but socially excessive. Fourth, weaker first-stage screening should increase probationary hiring, early separation, and post-hire rematching by shifting selection toward the early employment relationship.

The paper proceeds as follows. This section's remainder discusses related literature. Section 2 presents the environment of worker type and score-based rank. Section 3 characterizes assignment, aggregate output, and competitive decentralization, before Section 4 compares pre-AI and AI regimes. Section 5 studies endogenous application AI, and Section 6 analyzes post-hire rematching and probationary hiring. Finally, Section 7 states four empirical predictions outlines designs for testing them before Section 8 concludes.

Related Literature. The paper contributes first to the literature on assignment, sorting, and matching. Becker (1973) and Sattinger (1993) show that many economic environments are fundamentally assignment problems in which output depends on who is matched with whom. Shimer and Smith (2000) study assortative matching in decentralized search markets. The distinctive feature here is that assignment operates on noisy public rank rather than worker type. That shift makes the informational content of first-stage screening

a primitive object and permits a clean separation between sorting on observables and type.

Second, the paper relates to classic work on signaling, employer learning, and post-hire selection. A sizable literature starting with Spence (1973) shows how costly signals support separating equilibria in labor markets. Jovanovic (1979), Altonji and Pierret (2001), and Lange (2007), among others, study how match quality is learned over time. Loh (1994) and Wang and Weiss (1998) analyze probation and early separation as sorting devices. The present analysis is complementary to those contributions. Rather than modeling a specific signaling or learning technology, it embeds imperfect pre-match screening, competitive matching, and post-hire correction inside a common framework.

Third, the paper connects to the growing literature on GenAI and productivity. Experimental and field evidence shows that AI can raise output, often with heterogeneous effects across workers and tasks (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Dell'Acqua, McFowland III, Mollick, Lifshitz, Kellogg, Rajendran, Krayer, Candelon, and Lakhani, 2026). That literature speaks to how AI changes performance of a worker assigned to a task. The contribution here is to show that within-task productivity gains need not result in aggregate labor-market gains when screening becomes less informative.

Finally, the paper contributes to the emerging literature on AI in hiring and personnel selection. Fernandez-Mateo (2025) provides a broad overview of how GenAI is reshaping hiring processes. Lievens and Dunlop (2025) discuss the implications of applicant use of GenAI for the validity of selection procedures. Hickman, Dunlop, and Wolf (2024) show that large language models already perform strongly on some high-stakes verbal tests, raising concerns for unproctored screening. Canagasuriam and Lukacik (2025) study AI-enabled cheating in asynchronous video interviews. Robie, Wingate, Baytalskaya, and Butera (2026) document substantial self-reported use of GenAI in pre-hire assessments and show that warnings reduce such use. Particularly close in spirit are Galdin and Silbert (2025), who study how large language models lower the cost of written labor-market signaling, Cowgill, Hernandez-Lagos, and Wright (2026), analyzing how applicant access to GenAI can alter screening accuracy, and Cui, Dias, and Ye (2025), who show that AI-generated cover letters can increase tailoring and callback rates while weakening the relationship between application tailoring and hiring outcomes. The paper complements empirical work with a stylized general-equilibrium theory of sorting, wages, endogenous applicant AI use, and post-hire organizational adjustment.

2 Model

The model focuses on three primitives through which AI affects labor-market outcomes: direct productivity, the cost of mismatch within a realized worker-firm pair and the screening technology that maps latent type into observable pre-match scores.

2.1 Environment

There is a unit mass of workers and a unit mass of firms. Worker and firm latent types are denoted by X and Y, respectively, and both are uniformly distributed on $[0,1]$. Worker type x is a one-dimensional latent fit index. It bundles two forces that are distinct but intentionally summarized here on a common line: a vertical component that affects productivity at every firm and a horizontal component that governs where along the task space the worker is best deployed.[2] Firm type y indexes the task location the firm would most like to fill. Under full information, the efficient partner of firm y is therefore worker $x = y$.

In regime τ, $\tau \in \{0 := Pre - AI, 1 := AI\}$, output produced by worker x at firm y is

$$\pi_\tau(x,y) = q_\tau(x) - C_\tau(x-y)^2, \tag{1}$$

where $C_\tau > 0$ and $q_\tau : [0,1] \to \mathbb{R}$ is bounded, weakly increasing, and absolutely continuous. The function q_τ captures worker-specific productivity independent of firm type. The quadratic term captures the loss from assigning the worker to the wrong task location. Positive assortative matching on types is therefore efficient under full information.

The support and marginal distributions are normalizations for the assignment logic. Our main qualitative results depend only on the order structure induced by types and ranks.

2.2 Screening Scores and Market Rank

Before matching takes place under regime τ, each worker generates an observable screening score. This score may summarize an application file, a test outcome, a portfolio, or

[2] This bundling is a tractability device common in assignment models. The object here is not multidimensional skill heterogeneity per se, but the effect of noisy screening on allocation when a single latent fit dimension is decisive for efficient assignment (Sattinger, 1993).

a short work sample. In regime τ, a worker of type x generates score S_τ from a regime-specific screening technology. Let $g_\tau(s \mid x)$ and $G_\tau(s \mid x)$ denote the conditional density and distribution of $S_\tau \in \mathcal{S}_\tau \subset \mathbb{R}$ given $X = x$.

Because matching is one-to-one, the allocation depends on a worker's position in the cross-sectional score distribution rather than on the score's cardinal units. Let

$$H_\tau(s) := \Pr(S_\tau \leq s)$$

denote the distribution function of scores, and define the worker's (visible) market rank by

$$R_\tau := H_\tau(S_\tau).$$

Then R_τ is uniform on $[0, 1]$.

Assumption 1 (Screening technology)**.** For each regime $\tau \in \{0, 1\}$:

(i) $\mathcal{S}_\tau$ is an interval, the conditional density $g_\tau(s \mid x)$ is jointly continuous and strictly positive on $\mathcal{S}_\tau \times [0, 1]$, and H_τ is continuous and strictly increasing on $\mathcal{S}_\tau$;

(ii) for every $x' > x$, the likelihood ratio $s \mapsto g_\tau(s \mid x')/g_\tau(s \mid x)$ is strictly increasing;

(iii) for almost every $s \in \mathcal{S}_\tau$, the posterior distribution of X conditional on $S_\tau = s$ is nondegenerate.

Assumption 1 imposes a transparent structure on first-stage screening. Higher worker types shift the score distribution upward in the monotone-likelihood-ratio sense, but condition (iii) rules out perfect revelation. Screening is informative in each regime but never perfect.

A useful interpretation distinguishes common separable polish from genuine garbling. A common level shift in scores, or any strictly increasing transformation of scores, changes appearance without changing market rank and is therefore allocatively irrelevant.[3] By contrast, score compression—higher types gaining less in score than lower types—or additional type-blind noise increases overlap in the score distributions of different worker types and weakens the mapping from latent type into observed rank. The dependence-order

[3] Accordingly, the relevant empirical question is not whether AI raises raw application scores or improves their tone, fluency, or format on average. It is whether AI changes the induced ordering or changes what that ordering reveals about latent fit.

comparisons used below are intended to capture these economically distinct channels in reduced form. The Gaussian specification nests them through changes in the loading of type on score and, more generally, through the signal-to-noise ratio.

Because R_τ is a strictly increasing transformation of S_τ, conditioning on score and conditioning on market rank are equivalent. For each regime define the posterior objects

$$m_\tau(r) := \mathbb{E}[X \mid R_\tau = r], \qquad v_\tau(r) := \mathrm{Var}(X \mid R_\tau = r) \quad \text{and} \quad \bar{q}_\tau(r) := \mathbb{E}[q_\tau(X) \mid R_\tau = r].$$

The posterior mean $m_\tau(r)$ is the expected worker type at market rank r, $v_\tau(r)$ is the residual within-rank uncertainty and $\bar{q}_\tau(r)$ is expected worker productivity conditional on rank.

A convenient derived summary of screening quality is the rank-correlation statistic

$$\rho_\tau := 12\mathbb{E}[XR_\tau] - 3. \tag{2}$$

Since both X and R_τ are uniform on $[0,1]$, ρ_τ is Spearman's rank correlation between worker type and observed screening rank.

Lemma 1 (Informative but imperfect screening). *Under Assumption 1, the map m_τ is continuous and strictly increasing on $[0,1]$, and the derived rank-correlation statistic satisfies*

$$0 < \rho_\tau < 1.$$

The statistic ρ_τ is a reduced-form summary of the underlying score technology. It summarizes the allocative content of first-stage screening.

Conditional on rank, expected output at firm y is

$$\begin{aligned} \bar{\pi}_\tau(r, y) &:= \mathbb{E}[\pi_\tau(X, y) \mid R_\tau = r] \\ &= \bar{q}_\tau(r) - C_\tau\Big(v_\tau(r) + \big(m_\tau(r) - y\big)^2\Big). \end{aligned} \tag{3}$$

The economy therefore becomes a matching problem on ranks. Firms sort workers on the basis of the posterior means encoded by those ranks while bearing the residual within-rank uncertainty that screening fails to eliminate.

3 Analysis of the Baseline Economy

A social planner first chooses the best assignment given ranks. The resulting allocation is then decentralized as an assignment game. The Gaussian specification at the end of the section provides an illustrative example.

3.1 Assignment and Aggregate Output

An assignment rule is a measure-preserving map $\gamma : [0,1] \to [0,1]$. Measure preservation reflects one-to-one matching, i.e., the mass of workers assigned to any set of firms equals the mass of those firms.

In regime τ, $\tau \in \{0,1\}$, the planner chooses γ to solve

$$\max_{\gamma} \int_0^1 \bar{\pi}_\tau\big(r, \gamma(r)\big)\, dr.$$

Proposition 1 (Sorting on visible ranks). *Under Assumption 1, the unique surplus-maximizing assignment in regime τ is*

$$\gamma_\tau(r) = r \qquad \text{for all } r \in [0,1].$$

The allocation is therefore efficient conditional on observables. The market does not fail by sorting incorrectly on rank, but it fails because visible rank need not coincide with worker type. The inefficiency is informational rather than conditional on what the market observes.

Define the average worker-specific productivity component by

$$Q_\tau := \int_0^1 q_\tau(x)\, dx.$$

Theorem 1 (Aggregate output and the informational wedge). *Aggregate output in regime τ is*

$$Y_\tau = Q_\tau - \frac{C_\tau}{6}(1 - \rho_\tau). \tag{4}$$

Under full information,

$$Y_\tau^{\mathrm{FI}} = Q_\tau. \tag{5}$$

Hence the allocative loss generated by imperfect first-stage screening is

$$\Delta_\tau := Y_\tau^{\mathrm{FI}} - Y_\tau = \frac{C_\tau}{6}(1 - \rho_\tau). \tag{6}$$

Theorem 1 shows that the allocative role of first-stage screening collapses to a single wedge. The term Q_τ captures the productive value created inside realized matches. The term $C_\tau(1 - \rho_\tau)/6$ captures the output lost because workers are allocated on the basis of a noisy ranking rather than fit. Mismatch costs scale the harm created by a given informational error, and tighter alignment between type and rank attenuates it. The model therefore separates technological gains inside matches from informational losses in the assignment process.

3.2 Competitive Decentralization

Because utility is transferable, the planner's allocation can be decentralized as an assignment game. A competitive equilibrium in regime τ is a triple $(w_\tau, J_\tau, \gamma_\tau)$ consisting of a worker payoff schedule $w_\tau : [0,1] \to \mathbb{R}$, a firm profit schedule $J_\tau : [0,1] \to \mathbb{R}$, and a measure-preserving assignment rule γ_τ such that

$$w_\tau(r) + J_\tau(y) \geq \bar{\pi}_\tau(r, y) \qquad \text{for all } (r, y) \in [0,1]^2, \tag{7}$$

with equality whenever $y = \gamma_\tau(r)$.

Proposition 2 (Competitive equilibrium and representative normalization). *For each regime τ, $\tau \in \{0,1\}$, a competitive equilibrium exists and every equilibrium implements the assignment $\gamma_\tau(r) = r$ from Proposition 1. Any supporting worker payoff schedule is absolutely continuous and satisfies*

$$w_\tau(r) = \underline{w}_\tau + \int_0^r \partial_1 \bar{\pi}_\tau(t, t)\, dt \tag{8}$$

for some constant $\underline{w}_\tau \in \mathbb{R}$. Hence equilibrium payoffs are unique up to an additive constant.

Among the supporting equilibria, the firm-optimal representative equilibrium is characterized by

$$w_\tau(0) = 0. \tag{9}$$

In that representative equilibrium,

$$w_\tau(r) = \int_0^r \partial_1 \bar{\pi}_\tau(t,t)\, dt \quad \textit{and} \quad J_\tau(y) = \bar{\pi}_\tau(y,y) - w_\tau(y). \tag{10}$$

The allocation is pinned down uniquely, but the transfer schedule is not. That indeterminacy is standard in assignment games. The firm-optimal representative equilibrium is adopted because the later decentralized extensions depend on firms' continuation profits. Under this normalization the lowest rank receives zero rent, and firms obtain the largest continuation value consistent with support of the efficient visible-rank assignment.

Corollary 1 (**Full-information benchmark**). *Under full information the efficient assignment is $x = y$. In the firm-optimal representative equilibrium,*

$$w_\tau^{\mathrm{FI}}(x) = q_\tau(x) - q_\tau(0) \quad \textit{and} \quad J_\tau^{\mathrm{FI}}(y) = q_\tau(0).$$

Under full information all worker-specific productivity rents are bid into wages under the chosen normalization, so the representative firm's continuation value is constant across firm types. That property makes the post-hire outside option tractable in the analysis below.

3.3 Gaussian Illustration

The Gaussian specification provides a cardinal representation of the same rank-based environment. Let worker and firm types be standard normal, $X, Y \sim N(0,1)$. In regime τ suppose output is

$$\pi_\tau(x,y) = \eta_\tau + \theta_\tau x - C_\tau (x-y)^2,$$

where $\eta_\tau \in \mathbb{R}$ and $\theta_\tau > 0$. Screening is generated by the score

$$S_\tau = \mu_\tau + \beta_\tau X + \varepsilon_\tau, \qquad \varepsilon_\tau \sim N(0, \sigma_\tau^2),$$

with $\beta_\tau > 0$ and ε_τ independent of X. Define

$$D_\tau := \beta_\tau^2 + \sigma_\tau^2 \quad \text{and} \quad \kappa_\tau := \frac{\beta_\tau}{\sqrt{D_\tau}} \in (0,1).$$

κ_τ is the cardinal counterpart of screening informativeness. A larger κ_τ means that the screening score is more tightly aligned with worker type.

Proposition 3 (Gaussian illustration). *In the Gaussian environment:*

(i) the worker with score realization s is matched to firm type

$$y_\tau(s) = \frac{s - \mu_\tau}{\sqrt{D_\tau}};$$

(ii) the posterior distribution of worker type satisfies

$$\mathbb{E}[X \mid S_\tau = s] = \frac{\beta_\tau}{D_\tau}(s - \mu_\tau) = \kappa_\tau y_\tau(s) \quad \textit{and} \quad \mathrm{Var}(X \mid S_\tau = s) = 1 - \kappa_\tau^2;$$

(iii) aggregate output is

$$Y_\tau = \eta_\tau - 2C_\tau(1 - \kappa_\tau); \tag{11}$$

(iv) in the firm-optimal representative equilibrium, the wage schedule is

$$w_\tau(y) = \theta_\tau \kappa_\tau y + C_\tau \kappa_\tau (1 - \kappa_\tau) y^2. \tag{12}$$

The Gaussian formulas isolate the same structure in closed form. The level shift μ_τ has no allocative consequence because only ordinal position matters. The informativeness parameter κ_τ governs the quality of sorting. The slope parameter θ_τ governs the reward attached to visible rank without changing the assignment rule itself. Better-looking score levels and better screening are therefore distinct objects.

4 Comparing the Pre-AI and AI Regimes

This section isolates the paper's central mechanism. AI is allowed to improve what happens inside every realized match and to reduce mismatch losses within a match. Any decline in aggregate output must therefore pass through the informational content of first-stage screening outputs.

Assumption 2 (AI comparison on primitives). The two regimes satisfy:

(i) $q_1(x) \geq q_0(x)$ for all $x \in [0, 1]$, with strict inequality on a set of positive measure;

(ii) the gain $q_1 - q_0$ is weakly increasing on $[0,1]$;

(iii) $0 < C_1 \leq C_0$;

(iv) for every bounded increasing pair of functions $\phi, \psi : [0,1] \to \mathbb{R}$,

$$\mathbb{E}\big[\phi(X)\psi(R_1)\big] \leq \mathbb{E}\big[\phi(X)\psi(R_0)\big].$$

Assumption 2 (iv) requires the AI-era screening stage to generate a weaker positive association between type and observed market position. A sufficient case is an order-preserving garbling of pre-AI screening scores before final rank is formed.[4]

Proposition 4 (AI weakens rank alignment). *Under Assumption 2 (iv),*

$$\rho_1 \leq \rho_0. \tag{13}$$

The dependence order therefore maps directly into the summary statistic that enters aggregate output.

Proposition 5 (Pairwise productivity dominance). *Under Assumption 2,*

$$\pi_1(x,y) \geq \pi_0(x,y) \qquad \text{for all } (x,y) \in [0,1]^2,$$

with strict inequality on a set of positive measure.

Proposition 5 isolates the direct technological effect of AI. Holding the realized match fixed, every worker-firm pair produces at least as much under AI as before, and typically more. Any aggregate decline must therefore come from who is matched to whom rather than from what a fixed match can produce.

Theorem 2 (Higher pairwise productivity, lower aggregate output). *Under Assumption 2, aggregate output satisfies*

$$Y_1 - Y_0 = (Q_1 - Q_0) - \frac{1}{6}\Big(C_1(1-\rho_1) - C_0(1-\rho_0)\Big). \tag{14}$$

[4] Economically, the comparison covers two familiar channels: compression of the type-dependent component of screening outputs and the addition of type-blind noise. Pure monotone polishing is excluded because Proposition 6 makes it allocatively irrelevant.

Hence

$$Y_1 < Y_0 \quad \Longleftrightarrow \quad Q_1 - Q_0 < \frac{1}{6}\Big(C_1(1-\rho_1) - C_0(1-\rho_0)\Big). \tag{15}$$

In particular, if $Y_1 < Y_0$, then necessarily

$$C_1(1-\rho_1) > C_0(1-\rho_0). \tag{16}$$

Theorem 2 formalizes the central insight. AI can raise the productivity of every worker-firm pair and still lower aggregate output because the market allocates workers on the basis of a noisier proxy for worker type. The role of $C_1 \leq C_0$ is important. A lower mismatch cost works against the possibility of aggregate decline because it makes any given mismatch less damaging once it occurs. The result is therefore conservative: aggregate output falls despite the fact that AI both raises direct productivity and softens the loss from mismatch inside realized pairs. In other words, every worker becomes more productive at any firm. A negative aggregate effect requires a sufficiently large deterioration in first-stage information.

When $q_1 - q_0$ is increasing, AI also magnifies direct productivity differences in the upper part of the worker distribution. The return to climbing the visible ranking then becomes steeper, even though the assignment rule itself is unchanged. This force matters when workers can choose how intensively to use AI in producing screening outputs in Section 5 below.

Proposition 6 (Ordinal irrelevance of polishing). *Fix production primitives (q, C). Suppose two screening-score procedures generate the same observed rank almost surely. Then they induce the same posterior objects, the same assignment, the same aggregate output, and the same competitive payoffs up to the usual additive normalization of wages.*

Proposition 6 draws a sharp line between polishing and informational change. A transformation of screening outputs that leaves ranks unchanged is allocatively irrelevant in the baseline economy. What matters is whether AI changes rank or changes what rank order reveals about worker type.

Gaussian Illustration. In the Gaussian illustration, the primitive comparison reduces to $\kappa_1 \leq \kappa_0$, together with $\eta_1 + \theta_1 x \geq \eta_0 + \theta_0 x$ for all x and $C_1 \leq C_0$. The output comparison

then becomes

$$Y_1 - Y_0 = (\eta_1 - \eta_0) - 2\Big(C_1(1-\kappa_1) - C_0(1-\kappa_0)\Big).$$

The Gaussian representation therefore mirrors the general logic exactly. Higher pairwise productivity can coexist with lower aggregate output when the degradation in first-stage information is sufficiently severe.

5 Endogenous Application AI

This section isolates applicant-side AI use as a strategic investment in rank. The central question is whether decentralized incentives to improve application materials align with social value once all workers can adjust simultaneously.

Fix production primitives (q, C) and let $A := [0, \bar{a}]$ denote the set of application-AI intensities available to workers. When all workers choose the same intensity $a \in A$, each worker submits an application package summarized by an observable score $S(a)$ with atomless distribution function F_a, and market rank is

$$R_a := F_a\big(S(a)\big).$$

The induced symmetric screening technology therefore fits directly into the baseline model.

A worker who deviates to intensity b while the rest of the market uses a generates score $S(b)$ and is assigned the market position

$$R(b; a) := F_a\big(S(b)\big).$$

Application-AI use has cost $k : A \to \mathbb{R}_+$, where k is continuous and weakly increasing.

Assumption 3 (Application-AI score technology). The family $\{S(a) : a \in A\}$ satisfies:

(i) for every $a \in A$, the induced score $S(a)$ and the induced market rank $R_a := F_a(S(a))$ satisfy Assumption 1;

(ii) for every $a' > a$, and every bounded increasing pair $\phi, \psi : [0,1] \to \mathbb{R}$,

$$\mathbb{E}\big[\phi(X)\psi(R_{a'})\big] \le \mathbb{E}\big[\phi(X)\psi(R_a)\big];$$

(iii) application AI affects only the screening stage and not production inside a realized match.

Assumption 3 captures symmetric applicant AI that is primarily mimicry or garbling rather than pure level improvement. Common use may raise the polish of all applications, but it either compresses the type-dependent component of scores or adds type-blind noise, thereby weakening the mapping from type into market rank. By contrast, a common improvement that leaves the induced rank unchanged falls under Proposition 6 and is allocatively irrelevant in the baseline economy.

Define the derived screening-quality statistic under symmetric use a by

$$\rho(a) := 12\mathbb{E}[XR_a] - 3. \tag{17}$$

Let w_a denote the firm-optimal representative-equilibrium wage schedule induced by the symmetric screening technology at a.

Lemma 2 (Application AI and screening quality). *Under Assumption 3, the map $a \mapsto \rho(a)$ is weakly decreasing on A.*

A worker who deviates from the market profile a to intensity b obtains private payoff

$$U(b; a) := \mathbb{E}\big[w_a(R(b; a))\big] - k(b).$$

Aggregate welfare under the symmetric profile a is

$$W(a) := Q - \frac{C}{6}\big(1 - \rho(a)\big) - k(a), \tag{18}$$

where $Q := \int_0^1 q(x)\,dx$.

Proposition 7 (Efficient benchmark). *Under Assumption 3, the profile*

$$a^* = 0$$

is welfare maximizing.

In this environment application AI is purely positional. Symmetric increases in its use do not improve productive capacity inside matches. They only degrade the information

embedded in first-stage screening and absorb resources. The efficient benchmark therefore sets application-AI intensity at its lowest feasible level.

Theorem 3 (Excessive application AI). *Suppose U is continuous on $A \times A$ and, for each $a \in A$, the map $b \mapsto U(b; a)$ is continuously differentiable and strictly concave. Define the best-response map*

$$B(a) := \arg\max_{b \in A} U(b; a).$$

If

$$\partial_b U(0; 0) > 0 \qquad \text{and} \qquad \partial_b U(\bar{a}; \bar{a}) < 0,$$

then there exists an interior symmetric equilibrium $a^e \in (0, \bar{a})$ such that $a^e = B(a^e)$. Every such equilibrium satisfies

$$a^e > a^*.$$

The wedge is strategic rather than technological.[5] A worker privately values application AI because it can improve market rank against a given wage schedule. The interaction therefore has a prisoner's-dilemma structure. When all workers raise application-AI intensity, productive matches do not improve, the screening stage becomes noisier, and resource costs rise.

The force behind excessive use is stronger when direct productivity is steeper at the top. If AI raises the slope of q in the upper tail, the payoff gradient attached to visible rank becomes sharper. The private incentive to manipulate screening outputs then rises even though the social value of common polish does not.

5.1 Gaussian Illustration

Take the Gaussian environment from Proposition 3 with fixed production primitives (η, θ, C). Let application AI affect screening according to

$$S(X, a) = \mu(a) + \beta(a)X + \varepsilon, \qquad \varepsilon \sim N(0, \sigma^2),$$

[5] Application AI does not raise output inside the job in this section. Its private value is purely positional, i.e., a worker wants a higher place in the visible ranking given the market wage schedule.

where $\mu'(a) \geq 0$ and $\beta'(a) \leq 0$. Here $\mu(a)$ captures pure polish or a common level improvement, whereas $\beta(a)$ captures how strongly latent type loads into the score. The Gaussian illustration therefore makes explicit the distinction between allocatively irrelevant level shifts and harmful compression. Define

$$D(a) := \beta(a)^2 + \sigma^2 \quad \text{and} \quad \kappa(a) := \frac{\beta(a)}{\sqrt{D(a)}}.$$

Then $\kappa(a)$ is weakly decreasing in a, which is the Gaussian counterpart of Assumption 3(ii).

If the market profile is a, a worker who deviates to b is assigned the market position

$$y_{b|a} = \frac{\mu(b) - \mu(a)}{\sqrt{D(a)}} + \frac{\beta(b)}{\sqrt{D(a)}} X + \frac{\varepsilon_b}{\sqrt{D(a)}},$$

where ε_b is an independent copy of ε, so that $y_{b|a} \sim N(\alpha_{b|a}, \delta^2_{b|a})$ with

$$\alpha_{b|a} := \frac{\mu(b) - \mu(a)}{\sqrt{D(a)}} \quad \text{and} \quad \delta^2_{b|a} := \frac{\beta(b)^2 + \sigma^2}{D(a)}.$$

Using the wage schedule in Equation (12), private payoff is

$$U(b; a) = \theta\kappa(a)\alpha_{b|a} + C\kappa(a)\big(1 - \kappa(a)\big)\big(\alpha^2_{b|a} + \delta^2_{b|a}\big) - k(b),$$

while social welfare under symmetric use is

$$W(a) = \eta - 2C\big(1 - \kappa(a)\big) - k(a).$$

The Gaussian illustration separates polish from information in closed form. The level shift $\mu(a)$ affects private incentives through relative market position, but it drops out of social welfare. By contrast, a lower loading $\beta(a)$ reduces $\kappa(a)$ and directly worsens screening quality.

6 Selection Inside the Firm

When first-stage screening is less informative, part of the selection problem is shifted from the market to the firm. Two margins are immediate. Firms may separate and reenter the market after learning more about incumbent workers, and they may write probationary contracts that preserve the right to terminate after a short evaluation period.

6.1 Costly Rematching after Employer Learning

Consider a two-period environment. In period 1, the baseline competitive equilibrium from Section 3.2 forms matches and wages on the basis of first-stage ranks. At the end of period 1, the first-period employer privately learns the worker's type X. The employer then decides whether to retain the worker for period 2 at the inherited wage or to separate and enter a secondary market.

Throughout this subsection, the period-1 wage schedule is inherited into period 2 for any retained worker. After the incumbent employer learns X at the end of period 1, the employer alone chooses whether to continue the relationship at that inherited wage or to separate and enter the secondary market. Only workers who are separated at that stage participate in the secondary market. Retained workers do not renegotiate or move before period-2 output is realized.[6] The extension therefore isolates employer-side correction of an initially noisy match.

Fix a baseline regime and its firm-optimal representative-equilibrium wage schedule w. A firm of type y is initially matched with the worker of visible rank y. If the firm retains its incumbent after learning latent type x, continuation profit is

$$\Pi(x, y) := q(x) - C(x - y)^2 - w(y). \tag{19}$$

If the firm separates, it pays a rematching cost $k_R > 0$ and enters a secondary market for separated experienced workers. The maintained reduced-form assumption is that realized performance is sufficiently informative in that market that types are publicly

[6] This restriction isolates employer-side correction of noisy initial matches. Allowing bilateral renegotiation would affect the division of second-period surplus, but not the basic force emphasized here. When first-stage screening is weaker, the option to revisit the match after richer information arrives becomes more valuable.

observed among participants. By Corollary 1, the firm-optimal continuation value of entering that market is then

$$\Gamma := q(0) - k_R. \tag{20}$$

The outside option is constant across incumbent firms because, under the representative normalization, all worker-specific rents are bid into wages in the experienced-worker market.

Proposition 8 (Post-hire rematching and option value). *A firm of type y retains its worker if and only if*

$$\Pi(x, y) \geq \Gamma. \tag{21}$$

Its expected continuation value is

$$\mathcal{J}(y) = \mathbb{E}\big[\max\{\Pi(X, y), \Gamma\} \mid R = y\big], \tag{22}$$

and the option value of access to the secondary market is

$$\Omega(y) := \mathcal{J}(y) - \mathbb{E}\big[\Pi(X, y) \mid R = y\big] = \mathbb{E}\big[(\Gamma - \Pi(X, y))_+ \mid R = y\big]. \tag{23}$$

If q is twice continuously differentiable and $q''(x) \leq 2C$ for all $x \in [0, 1]$, then $x \mapsto \Pi(x, y)$ is concave for every y, so the retention set is an interval in latent type.

Under the maintained timing, Proposition 8 is the incumbent firm's continuation problem at an inherited wage. The object $\Omega(y)$ is the value of correcting a match that looks attractive ex ante but proves weak once the employer learns more. It is therefore a direct measure of how fragile first-stage matching is for a firm of type y. Better initial screening shrinks $\Omega(y)$ because fewer apparently favorable matches unravel after worker type is revealed.

Corollary 2 (Weaker screening and rematching). *Consider two first-stage screening technologies, indexed by $i \in \{0, 1\}$, that induce the same outside option Γ. Fix a firm type y. If the conditional continuation-profit distribution*

$$\Pi_i(X, y) \mid R_i = y$$

is a mean-preserving spread under $i = 1$ of the corresponding distribution under $i = 0$, then

$$\Omega_1(y) \geq \Omega_0(y). \tag{24}$$

If the spread is strict on a set of firm types of positive measure, average option value rises strictly.

The corollary isolates the relevant mechanism. What makes rematching valuable is not a lower mean worker payoff by itself. It is a thicker lower tail of post-learning continuation values. A noisier first-stage screen creates more hidden mismatch risk inside apparently attractive period-1 matches, and that is exactly the risk a secondary market helps firms undo.

6.2 Probationary Hiring

Suppose the rematching option in Section 6.1 can be purchased ex ante. Suppose that at the hiring stage a firm may pay a cost $k_P > 0$ to make the employment relationship probationary. A probationary contract preserves the right to separate after the employer learns the worker's type at the end of period 1. Without probation, the match becomes permanent for period 2.

Proposition 9 (Probationary contracting)**.** *For a match formed at firm type y, the value of a permanent contract is*

$$J^N(y) = \mathbb{E}\big[\Pi(X, y) \mid R = y\big], \tag{25}$$

where $\Pi(X, y)$ is defined in Equation (19). The value of a probationary contract is

$$J^P(y) = \mathbb{E}\big[\max\{\Pi(X, y), \Gamma\} \mid R = y\big] - k_P = J^N(y) + \Omega(y) - k_P. \tag{26}$$

Hence a firm chooses probation if and only if

$$\Omega(y) \geq k_P. \tag{27}$$

Probationary hiring is the ex ante contractual counterpart to ex post rematching. A firm pays for probation precisely when first-stage screening leaves enough downside risk inside an ostensibly favorable match that the option to terminate later is worth purchasing in advance.

Corollary 3 (Weaker screening and probationary hiring). *Consider two first-stage screening technologies, indexed by $i \in \{0,1\}$, with the same outside option Γ and the same probation cost k_P. If, for every firm type y, the distribution $\Pi_i(X,y) \mid R_i = y$ under $i = 1$ is a mean-preserving spread of the corresponding distribution under $i = 0$, then every firm that chooses probation under $i = 0$ also chooses probation under $i = 1$. If the spread is strict on a set of positive measure, the mass of probationary contracts rises weakly and strictly whenever the threshold in Equation (27) is crossed on a set of positive measure.*

Once first-stage screening becomes less informative, part of the selection problem shifts from ex ante screening toward ex post selection inside the firm. Costly rematching and probationary hiring are two margins through which that shift appears.

6.3 Gaussian Illustration

Maintaining the Gaussian environment from Proposition 3, continuation profit from retaining the incumbent worker is

$$\Pi(x,y) = \eta + \theta x - C(x-y)^2 - \theta\kappa y - C\kappa(1-\kappa)y^2,$$

while the value of entering the secondary market is $\Gamma = \eta - k_R$. Retention therefore requires

$$\theta x - C(x-y)^2 \geq \theta\kappa y + C\kappa(1-\kappa)y^2 - k_R.$$

If the discriminant

$$D_R(y) := \theta^2 + 4C\Big(k_R + \theta(1-\kappa)y - C\kappa(1-\kappa)y^2\Big)$$

is nonnegative, the retention set is the interval $[x^-(y), x^+(y)]$, where

$$x^{\pm}(y) = \frac{2Cy + \theta \pm \sqrt{D_R(y)}}{2C}.$$

Since $X \mid y \sim N(\kappa y, 1-\kappa^2)$,

$$\Pr(\text{separation} \mid y) = \Phi\left(\frac{x^-(y) - \kappa y}{\sqrt{1-\kappa^2}}\right) + 1 - \Phi\left(\frac{x^+(y) - \kappa y}{\sqrt{1-\kappa^2}}\right).$$

The Gaussian continuation therefore delivers an explicit separation region for each incumbent firm type.

The corresponding probationary value is

$$\begin{aligned}\Omega(y) &= \mathbb{E}\big[(\eta - k_R - \Pi(X,y))_+ \mid y\big] \\ &= \int_{-\infty}^{x^-(y)} \big(\eta - k_R - \Pi(x,y)\big)\, \phi_{\kappa y,\, 1-\kappa^2}(x)\, dx \\ &\quad + \int_{x^+(y)}^{\infty} \big(\eta - k_R - \Pi(x,y)\big)\, \phi_{\kappa y,\, 1-\kappa^2}(x)\, dx,\end{aligned}$$

where $\phi_{m,v}$ denotes the Gaussian density with mean m and variance v. The firm chooses probation whenever $\Omega(y) \geq k_P$. Lower κ widens the tails of the conditional incumbent distribution and therefore raises both separation and probationary demand.

7 Predictions and Empirical Design

The discussion organizes the empirical content of the analysis into four direct predictions. Each is stated in terms of observable labor-market outcomes rather than model objects. Together they describe how applicant-facing AI should affect the predictive content of first-stage screening, the design of pre-hire assessment, the equilibrium use of application AI, and the importance of post-hire selection.

Prediction 1: AI makes first-stage screening less predictive of realized fit. The baseline mechanism does not require AI to lower average output inside tasks. It requires only that AI lower the informational content of the screening objects that determine visible rank. That implication is already consistent with several strands of early evidence. Cowgill, Hernandez-Lagos, and Wright (2026) show experimentally that access to GenAI can reduce screening accuracy in hiring and entrepreneurship settings. This interpretation is also consistent with Galdin and Silbert (2025), who show that once large language models lower the cost of producing customized written applications, employers' willingness to pay for that customization falls sharply and the resulting allocation becomes less meritocratic. Cui, Dias, and Ye (2025) find that an AI cover-letter tool increased textual tailoring and callback rates on Freelancer.com while substantially weakening the relationship between tailoring and callbacks. Lievens and Dunlop (2025) argue that applicant use of GenAI can

undermine the validity of selection procedures when firms rely on unproctored written materials. In the language of the model, observed scores may improve in level while the ordinal content of the induced ranking deteriorates. More broadly, recent evidence suggests that GenAI can flatten the quality distribution of produced artifacts by reducing the return to additional effort (Chen, Gong, Li, and Zhao, 2025). Although that setting is not a hiring environment, it is conceptually consistent with the concern that AI can compress informative variation in observable outputs.

Prediction 2: AI increases the use of harder-to-fake assessments and verification. Once first-stage documents become less informative, employers should shift weight toward selection methods that are harder to outsource to a language model or easier to verify directly.[7] The personnel-selection literature has long emphasized the predictive value of more behaviorally grounded screening instruments such as work-sample tests and structured assessments (Schmidt and Hunter, 1998; Roth, Bobko, and McFarland, 2005). Recent evidence suggests that this substitution is already under way. Hickman, Dunlop, and Wolf (2024) show that large language models perform strongly on some unproctored verbal ability tests, and Robie, Wingate, Baytalskaya, and Butera (2026) document both increasing candidate use of GenAI in high-stakes assessments and the mitigating effect of explicit warnings. Business reporting points in the same direction: firms have reintroduced in-person interviews, identity checks, and other verification devices in response to AI-assisted applications and interviews (Smith, 2025). The model interprets these responses as attempts to raise the effective informativeness of the screening process when low-friction artifacts cease to separate candidates well.

Prediction 3: Applicant-side AI use is privately attractive but socially excessive. In the model, application AI is privately attractive for a positional reason: it can improve visible rank against a given wage schedule. Symmetric adoption, however, does not improve productive capacity inside matches. It degrades the informativeness of screening and absorbs resources. The mechanism is therefore an arms race. Use should be most intense in markets where small improvements in visible position have large effects on expected payoffs, which is consistent with practitioner accounts of automated application behavior and AI-assisted tailoring (Ellis, 2024; The Economist, 2026). Evidence that GenAI raises

[7] This interpretation aligns with the personnel-selection evidence that behaviorally grounded assessments such as work-sample tests often dominate self-presented written materials in predictive validity (Schmidt and Hunter, 1998; Roth, Bobko, and McFarland, 2005).

productivity within tasks matters only indirectly here. It increases the value of obtaining a better assignment and thus steepens the private return to climbing the visible ranking (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025). The private diffusion of applicant AI should therefore exceed its social value in the screening stage.

Prediction 4. AI increases probationary hiring, early separation, and post-hire rematching. If firms cannot reliably identify fit before hiring, firms should do more learning after hiring. The post-hire section formalizes two margins: costly rematching and probationary contracting. Both imply more early-tenure separation when first-stage information deteriorates. This prediction is closely related to classic theories of job matching and employer learning (Jovanovic, 1979; Altonji and Pierret, 2001) and to earlier work on probation as a sorting device (Loh, 1994; Wang and Weiss, 1998). In current settings, the model suggests a specific mechanism: applicant-facing AI increases the downside tail of match quality among workers who initially look strong on observable materials. Firms should therefore rely more on probation, short trial projects, internship-like evaluation periods, and other contractual or organizational devices that preserve the option to revise an initially promising hire.

These predictions naturally do not imply that applicant-facing AI is always harmful, nor that every employer response will be efficient. The narrower implication is that the allocative consequences of AI depend on whether the technology raises or lowers the informativeness of the screening objects on which the market sorts. That perspective helps organize a rapidly evolving empirical domain. Increases in application volume, renewed interest in work-sample testing, greater use of verification, and higher value placed on post-hire evaluation need not be treated as disconnected phenomena. They are coherent responses to a common informational shock.

7.1 Measurement, Rank-Based Estimands, and Identification

The theory suggests a precise empirical test. The relevant estimand is not the Pearson correlation between raw application scores and later outcomes. The model is rank-based, so the appropriate objects are rank-order statistics and within-market predictive gradients: Spearman or Kendall associations between first-stage rank and later performance, the probability that a worker ranked above another worker ex ante also outperforms that

worker ex post, or the slope of validated performance on within-posting percentile rank.[8] A study that observes only callbacks, interview conversion, or raw first-stage scores cannot test the main mechanism, because those outcomes do not reveal whether the market is sorting on a more or less informative ordering.

One feasible design uses administrative hiring data from firms or platforms that observe applications, screening scores, hiring decisions, post-hire performance, and early exits. The relevant variation can come from staggered access to applicant AI tools, employer policies that restrict or permit AI use, randomized warnings or disclosure requirements of the kind studied by Robie, Wingate, Baytalskaya, and Butera (2026), or exposure differences across occupations in which applicant-facing AI is more or less useful. The central estimating question is whether the rank-order predictive content of pre-hire screening outputs falls when applicant-facing AI becomes more prevalent. The model allows average application quality or callback rates to rise at the same time.

Two comparative-static patterns are especially diagnostic. First, the decline in predictive content should be strongest where first-stage screening relies most heavily on artifacts that are easy to polish with AI—cover letters, writing samples, asynchronous interviews, and take-home submissions—and weaker where first-stage evaluation already depends on proctored tasks or live interaction. Second, conditional on the same percentile position in the first-stage ranking, later mismatch should rise more in those AI-exposed settings. These heterogeneity patterns help distinguish the mechanism from changes in applicant composition or broad labor-demand shocks, which can shift levels of hiring or performance without necessarily weakening the conditional relation between first-stage rank and later outcomes.

A complementary experimental design randomizes applicant access to an AI assistant when producing cover letters, writing samples, or responses to asynchronous interviews. Employers evaluate those materials in the usual way and, for a subset of applicants, richer validation measures are collected through proctored work-sample tests or standardized job tasks. The model predicts a particular sign pattern: AI access can raise first-stage evaluations and perhaps interview or callback rates while weakening the rank-order association between first-stage evaluations and validated performance. Pure score compression or

[8] This is why raw-score correlations can be misleading. Mechanical compression or ceiling effects can lower Pearson correlations even when the ordinal ranking is unchanged, whereas the model's allocative channel depends on the predictive content of the induced ordering.

ceiling effects are not sufficient. The relevant test is whether the induced ordering becomes less informative. Cowgill, Hernandez-Lagos, and Wright (2026) and Cui, Dias, and Ye (2025) show that designs of this general form are already empirically feasible.

The post-hire predictions can also be taken to the data. If applicant-facing AI weakens first-stage screening, firms more exposed to applicant AI should rely more on probationary hiring, work-sample verification, and early post-hire evaluation, and they should exhibit more early-tenure separation conditional on initial rank or initial screening percentile. These are not auxiliary implications. They are the organizational margins through which firms partially offset degraded public information.

8 Conclusion

Applicant-facing GenAI can improve applications and realized task performance while worsening the allocation of talent. Labor markets are organized around screening objects that are imperfect predictors of latent fit. The central implication is therefore not that AI mechanically disrupts sorting. It is that the market can continue to sort efficiently on a signal that has become less informative. Evidence that AI improves within-task productivity is therefore not sufficient to establish that it improves the allocation of talent.

The model studies this distinction in a one-to-one matching environment with transferable utility. Workers and firms have latent types, production depends on both worker-specific productivity and match quality, and the market observes only the visible ranks generated by noisy screening scores. The baseline allocation is efficient in the space of observables: workers are sorted to firms according to visible rank. The distortion arises because visible rank need not coincide with latent worker type. Aggregate output therefore differs from the full-information benchmark by a simple informational wedge governed by mismatch costs and by the alignment between latent type and observed rank.

This structure yields a sharp implication. AI can raise the productivity of every worker-firm pair and lower mismatch costs inside realized matches, yet aggregate output can still fall if first-stage screening becomes sufficiently less informative. The mechanism is not a collapse of sorting. It is sorting on a worse signal. A related implication is that purely cosmetic improvements in application materials have no allocative content unless they alter rank or change what rank reveals about latent type.

The decentralized and post-hire analyses translate the same mechanism into four empirical implications. First, applicant-facing AI should make first-stage screening less predictive of realized fit, even when applications look better on average. Second, firms should substitute toward harder-to-fake assessments and verification. Third, applicant-side AI use should be privately attractive but socially excessive. Fourth, weaker first-stage screening should increase probationary hiring, early separation, and post-hire rematching. Applicant-facing AI therefore affects not only hiring outcomes, but also wage schedules, turnover, and organizational design after hiring.

These predictions speak directly to an environment in which applications are increasingly AI-assisted and employers are experimenting with verification, work samples, and other higher-friction forms of assessment. A natural next step is empirical and experimental. Field designs that vary applicant access to GenAI together with employer screening technologies could trace the sequence emphasized in this paper: how AI changes screening outputs, how those changes affect assignment and wages, and how firms reorganize selection after hiring when pre-match information becomes less reliable.

References

Altonji, J. G., and C. R. Pierret (2001): "Employer Learning and Statistical Discrimination," *The Quarterly Journal of Economics*, 116(1), 313–350.

Becker, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813–846.

Brynjolfsson, E., D. Li, and L. R. Raymond (2025): "Generative AI at Work," *The Quarterly Journal of Economics*, 140(2), 889–942.

Canagasuriam, D., and E.-R. Lukacik (2025): "ChatGPT, Can You Take My Job Interview? Examining Artificial Intelligence Cheating in the Asynchronous Video Interview," *International Journal of Selection and Assessment*, Published online in 2025.

Chen, Y. J., J. Gong, J. Li, and Z. Zhao (2025): "Better Technology, Worse Motivation: GenAI's Mediocrity Trap," Hong Kong University, Working paper.

Cowgill, B., P. Hernandez-Lagos, and N. L. Wright (2026): "Does AI Cheapen Talk? Theory and Evidence From Global Entrepreneurship and Hiring," Discussion Paper 18442, IZA.

Cui, J., G. Dias, and J. Ye (2025): "Signaling in the Age of AI: Evidence from Cover Letters," Yale University, Working paper.

Dell'Acqua, F., E. McFowland III, E. Mollick, H. Lifshitz, K. C. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2026): "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," *Organization Science*, 37(2), 403–423.

Ellis, L. (2024): "You're Fighting AI With AI: Bots Are Breaking the Hiring Process," The Wall Street Journal, https://www.wsj.com/lifestyle/careers/ai-job-application-685f29f7, Accessed April 7, 2026.

Fernandez-Mateo, I. (2025): "How Generative AI Is Transforming Hiring in Organizations: Key Issues and Research Questions," *Journal of Organization Design*, Point of View article, published online September 11, 2025.

Galdin, A., and J. Silbert (2025): "Making Talk Cheap: Generative AI and Labor Market Signaling," Working paper, arXiv:2511.08785.

Hickman, L., P. D. Dunlop, and J. L. Wolf (2024): "The Performance of Large Language Models on Quantitative and Verbal Ability Tests: Initial Evidence and Implications for Unproctored High-Stakes Testing," *International Journal of Selection and Assessment*, 32, 499–511.

Jovanovic, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87(5), 972–990.

Lange, F. (2007): "The Speed of Employer Learning," *Journal of Labor Economics*, 25(1), 1–35.

Lievens, F., and P. D. Dunlop (2025): "Effects of Applicants' Use of Generative AI in Personnel Selection: Towards a More Nuanced View?," *International Journal of Selection and Assessment*, 33(1), e12516.

Loh, E. S. (1994): "Employment Probation as a Sorting Mechanism," *ILR Review*, 47(3), 471–486.

Noy, S., and W. Zhang (2023): "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381(6654), 187–192.

Robie, C., T. G. Wingate, N. Baytalskaya, and H. Butera (2026): "Candidate Generative AI Use in Pre-Hire Employment Assessments: Self-Reported Incidence and the Impact of Warnings," *International Journal of Selection and Assessment*, 34(2), e70056.

Roth, P. L., P. Bobko, and L. A. McFarland (2005): "A Meta-Analysis of Work Sample Test Validity: Updating and Integrating Some Classic Literature," *Personnel Psychology*, 58(4), 1009–1037.

Sattinger, M. (1993): "Assignment Models of the Distribution of Earnings," *Journal of Economic Literature*, 31(2), 831–880.

Schmidt, F. L., and J. E. Hunter (1998): "The Validity and Utility of Selection Methods in Personnel Psychology: Practical and Theoretical Implications of 85 Years of Research Findings," *Psychological Bulletin*, 124(2), 262–274.

Shimer, R., and L. Smith (2000): "Assortative Matching and Search," *Econometrica*, 68(2), 343–369.

Smith, R. A. (2025): "AI Is Forcing the Return of the In-Person Job Interview," The Wall Street Journal, https://www.wsj.com/lifestyle/careers/ai-job-interview-virtual-in-person-305f9fd0, Accessed April 7, 2026.

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

The Economist (2026): "Job Applicants Are Winning the AI Arms Race Against Recruiters," The Economist, https://www.economist.com/business/2026/01/12/job-applicants-are-winning-the-ai-arms-race-against-recruiters, Accessed April 7, 2026.

Wang, R., and A. Weiss (1998): "Probation, Layoffs, and Wage-Tenure Profiles: A Sorting Explanation," *Labour Economics*, 5(3), 359–383.

A Omitted Proofs

Proof of Lemma 1. Fix a regime τ and suppress the subscript. Because H is continuous and strictly increasing, conditioning on market rank $R = r$ is equivalent to conditioning on the score realization $S = H^{-1}(r)$. By Bayes' rule, the posterior density of X given $S = s$ is proportional to $g(s \mid x)$. Assumption 1(ii) implies that these posterior densities are ordered in the monotone-likelihood-ratio sense as s increases. It follows that the posterior distribution of X is strictly increasing in s, and therefore in r, in the sense of first-order stochastic dominance. Hence the posterior mean $m(r) = \mathbb{E}[X \mid R = r]$ is continuous and strictly increasing.

Since $m(R)$ is an increasing nonconstant function of the uniformly distributed variable R, Chebyshev's covariance inequality yields

$$\operatorname{Cov}(X, R) = \operatorname{Cov}(m(R), R) > 0,$$

so $\rho = 12 \operatorname{Cov}(X, R) > 0$. For the upper bound, Cauchy–Schwarz gives

$$\operatorname{Cov}(X, R) \le \sqrt{\operatorname{Var}(X) \operatorname{Var}(R)} = \frac{1}{12},$$

with equality only if $R = aX + b$ almost surely. Because both X and R are uniform on $[0, 1]$, equality would require $R = X$ almost surely, contradicting Assumption 1(iii). Therefore $\rho < 1$. □

Proof of Proposition 1. Take $r_1 < r_2$ and $y_1 < y_2$, and write $m_i := m_\tau(r_i)$. By Lemma 1, $m_1 < m_2$. Using Equation (3),

$$\begin{aligned}
&\left[\bar{\pi}_\tau(r_1, y_1) + \bar{\pi}_\tau(r_2, y_2)\right] - \left[\bar{\pi}_\tau(r_1, y_2) + \bar{\pi}_\tau(r_2, y_1)\right] \\
&= -C_\tau\left[(m_1 - y_1)^2 + (m_2 - y_2)^2 - (m_1 - y_2)^2 - (m_2 - y_1)^2\right] \\
&= 2C_\tau(m_2 - m_1)(y_2 - y_1) > 0.
\end{aligned}$$

Hence every inversion lowers total surplus. Repeated elimination of inversions yields the unique measure-preserving assignment $\gamma_\tau(r) = r$. □

Proof of Theorem 1. By Proposition 1, aggregate output is

$$Y_\tau = \mathbb{E}\left[\pi_\tau(X, R_\tau)\right] = \mathbb{E}[q_\tau(X)] - C_\tau \mathbb{E}\left[(X - R_\tau)^2\right].$$

Since X and R_τ are both uniform on $[0,1]$,

$$\mathbb{E}\big[(X - R_\tau)^2\big] = \frac{1}{12} + \frac{1}{12} - 2\,\mathrm{Cov}(X, R_\tau) = \frac{1-\rho_\tau}{6}.$$

Substituting yields Equation (4). Under full information, $R_\tau = X$ almost surely, so Equation (5) follows immediately, and hence so does Equation (6). □

Proof of Proposition 2. Fix regime τ and define

$$w_\tau(r) = \underline{w}_\tau + \int_0^r \partial_1 \bar{\pi}_\tau(t,t)\,dt, \qquad J_\tau(y) = \bar{\pi}_\tau(y,y) - w_\tau(y).$$

Condition Equation (7) is verified for the assignment $\gamma_\tau(r) = r$. If $r \le y$, then

$$\begin{aligned} w_\tau(r) + J_\tau(y) - \bar{\pi}_\tau(r,y) &= \bar{\pi}_\tau(y,y) - \bar{\pi}_\tau(r,y) - \int_r^y \partial_1 \bar{\pi}_\tau(t,t)\,dt \\ &= \int_r^y \Big(\partial_1 \bar{\pi}_\tau(t,y) - \partial_1 \bar{\pi}_\tau(t,t)\Big)dt \ge 0, \end{aligned}$$

where the inequality uses increasing differences of $\bar{\pi}_\tau$, established in the proof of Proposition 1. The case $r \ge y$ is analogous. Equality holds at $r = y$, so competitive equilibrium exists.

To establish uniqueness up to an additive constant, let $(\widetilde{w}_\tau, \widetilde{J}_\tau, \widetilde{\gamma}_\tau)$ be any equilibrium supporting the identity assignment. Equality on the diagonal yields

$$\widetilde{w}_\tau(r) + \widetilde{J}_\tau(r) = \bar{\pi}_\tau(r,r).$$

By the envelope theorem, $\widetilde{w}'_\tau(r) = \partial_1 \bar{\pi}_\tau(r,r)$ almost everywhere. Integrating proves Equation (8). The normalization Equation (9) then selects the representative equilibrium in Equation (10). □

Proof of Corollary 1. Under full information, $R_\tau = X$, $m_\tau(x) = x$, $v_\tau(x) = 0$, and $\bar{q}_\tau(x) = q_\tau(x)$. Applying Proposition 2 gives

$$w_\tau^{\mathrm{FI}}(x) = w_\tau^{\mathrm{FI}}(0) + \int_0^x q'_\tau(t)\,dt.$$

Under the normalization $w_\tau^{\mathrm{FI}}(0) = 0$, this is $q_\tau(x) - q_\tau(0)$. Firm profit is therefore

$$J_\tau^{\mathrm{FI}}(y) = q_\tau(y) - w_\tau^{\mathrm{FI}}(y) = q_\tau(0).$$

□

Proof of Proposition 3. Since $S_\tau \sim N(\mu_\tau, D_\tau)$, the worker with score realization s is matched to the firm at the same normal quantile, namely $y_\tau(s) = (s - \mu_\tau)/\sqrt{D_\tau}$. Standard normal updating yields

$$\mathbb{E}[X \mid S_\tau = s] = \frac{\beta_\tau}{D_\tau}(s - \mu_\tau) = \kappa_\tau y_\tau(s), \qquad \mathrm{Var}(X \mid S_\tau = s) = 1 - \kappa_\tau^2.$$

Moreover,

$$y_\tau(S_\tau) = \frac{\beta_\tau}{\sqrt{D_\tau}} X + \frac{\varepsilon_\tau}{\sqrt{D_\tau}},$$

so $\mathrm{Cov}(X, y_\tau(S_\tau)) = \kappa_\tau$ and both variables have variance 1. Hence

$$\mathbb{E}\left[(X - y_\tau(S_\tau))^2\right] = 2(1 - \kappa_\tau),$$

which gives Equation (11). Finally, conditional on assignment index y, expected surplus is

$$\bar{\pi}_\tau(y, \widetilde{y}) = \eta_\tau + \theta_\tau \kappa_\tau y - C_\tau\Big((1 - \kappa_\tau^2) + (\kappa_\tau y - \widetilde{y})^2\Big).$$

Applying Proposition 2 and the normalization $w_\tau(0) = 0$ gives

$$w_\tau'(y) = \partial_1 \bar{\pi}_\tau(y, y) = \theta_\tau \kappa_\tau + 2C_\tau \kappa_\tau (1 - \kappa_\tau) y,$$

which integrates to Equation (12). □

Proof of Proposition 4. Apply Assumption 2(iv) with $\phi(x) = x$ and $\psi(r) = r$. This gives

$$\mathbb{E}[XR_1] \le \mathbb{E}[XR_0].$$

Using the definition of ρ_τ in Equation (2) yields Equation (13). □

Proof of Proposition 5. By Equation (1),

$$\pi_1(x, y) - \pi_0(x, y) = \big(q_1(x) - q_0(x)\big) + (C_0 - C_1)(x - y)^2.$$

Under Assumption 2(i) and Assumption 2(iii), the right-hand side is nonnegative for every (x, y). It is strict on a set of positive measure because either $q_1(x) > q_0(x)$ on a positive-measure set or $C_1 < C_0$ and $x \neq y$. □

Proof of Theorem 2. Subtract Equation (4) across regimes. This yields Equation (14). The equivalence in Equation (15) is immediate. If $Y_1 < Y_0$, then the right-hand side of Equation (14) is negative. Since Assumption 2(i) implies $Q_1 - Q_0 \geq 0$, it follows that

$$C_1(1 - \rho_1) - C_0(1 - \rho_0) > 0,$$

which is exactly Equation (16). □

Proof of Proposition 6. If two screening-score procedures generate the same observed rank almost surely, then the conditional law of X given observed rank is the same under the two procedures. Hence the posterior objects coincide. The assignment and output statements then follow from Proposition 1 and Theorem 1. Competitive payoffs coincide up to an additive constant by Proposition 2. □

Proof of Lemma 2. Apply Assumption 3(ii) with $\phi(x) = x$ and $\psi(r) = r$. This yields

$$\mathbb{E}[XR_{a'}] \leq \mathbb{E}[XR_a] \qquad \text{whenever } a' > a.$$

Using the definition of $\rho(a)$ in Equation (17) gives the result. □

Proof of Proposition 7. Under Assumption 3(iii), symmetric welfare is given by Equation (18). By Lemma 2, the map $a \mapsto \rho(a)$ is weakly decreasing, and by assumption the cost function k is weakly increasing. Therefore $W(a)$ is weakly decreasing on A. It follows that $a = 0$ attains the maximum of W on A, so the profile $a^* = 0$ is welfare maximizing. □

Proof of Theorem 3. By continuity of U and strict concavity in the first argument, the best-response map B is single-valued and continuous. The condition $\partial_b U(0; 0) > 0$ implies $B(0) > 0$, while $\partial_b U(\bar{a}; \bar{a}) < 0$ implies $B(\bar{a}) < \bar{a}$. Hence the continuous function $B(a) - a$ is positive at $a = 0$ and negative at $a = \bar{a}$, so it has a zero in $(0, \bar{a})$. Any such fixed point is an interior symmetric equilibrium.

By Proposition 7, the profile $a^* = 0$ is welfare maximizing. Therefore any interior equilibrium a^e satisfies $a^e > a^*$. □

Proof of Proposition 8. After learning x, the firm compares the continuation value from retaining the incumbent, given by Equation (19), with the value of entering the secondary market, given by Equation (20). This yields the decision rule in Equation (21). Taking conditional expectations gives Equation (22). Subtracting expected retention profit yields Equation (23). Finally,

$$\frac{\partial^2 \Pi(x, y)}{\partial x^2} = q''(x) - 2C.$$

Hence $q''(x) \leq 2C$ implies that $x \mapsto \Pi(x, y)$ is concave for every y, so the superlevel set $\{x : \Pi(x, y) \geq \Gamma\}$ is an interval. □

Proof of Corollary 2. Write $Z_i(y) := \Pi_i(X, y) \mid R_i = y$ for $i \in \{0, 1\}$. By Equation (23),

$$\Omega_i(y) = \mathbb{E}\big[(\Gamma - Z_i(y))_+\big].$$

The function $z \mapsto (\Gamma - z)_+$ is convex. If $Z_1(y)$ is a mean-preserving spread of $Z_0(y)$, convex order implies

$$\mathbb{E}\big[(\Gamma - Z_1(y))_+\big] \geq \mathbb{E}\big[(\Gamma - Z_0(y))_+\big].$$

This is Equation (24). Strict inequality on a positive-measure set of firm types implies a strict increase in average option value. □

Proof of Proposition 9. Under a permanent contract, separation is unavailable, so the firm's expected continuation value is Equation (25). Under a probationary contract, the firm pays k_P and then chooses between retention and separation after learning X, which yields Equation (26). Comparing the two values gives the rule in Equation (27). □

Proof of Corollary 3. By Proposition 9, a firm chooses probation if and only if $\Omega(y) \geq k_P$. By Corollary 2, the maintained mean-preserving-spread condition implies $\Omega_1(y) \geq \Omega_0(y)$ for every y. Hence any firm that chooses probation under regime 0 also chooses probation under regime 1. Strict inequality on a positive-measure set yields the stated increase in the mass of probationary contracts whenever the threshold is crossed on a set of positive measure. □