

Marketing Science Institute Working Paper Series 2026

Report No. 26-129

Salesforce Compensation and Composition in the Age of AI

Saiquan Hu, Bo Pang, Gaoxiang Xu and Yuting Zhu

"Salesforce Compensation and Composition in the Age of AI" © 2026

Saiquan Hu, Bo Pang, Gaoxiang Xu and Yuting Zhu

MSI Working Papers are Distributed for the benefit of MSI corporate and academic members and the general public. Reports are not to be reproduced or published in any form or by any means, electronic or mechanical, without written permission.

Salesforce Compensation and Composition in the Age of AI

Saiquan Hu [*] Bo Pang [†] Gaoxiang Xu [‡] Yuting Zhu [§]

August 3, 2026

Abstract

Artificial intelligence is transforming how sales are produced, yet how salesforce compensation should adapt remains an open question. We develop a game-theoretic model in which a firm posts a linear compensation contract to attract and motivate sales candidates who differ in ability and outside opportunities. The contract operates along two interconnected margins: it induces effort from the hired salesperson and, through candidate self-selection, shapes the ability composition of the pool. Sales are jointly produced by human effort and AI through a constant-elasticity-of-substitution technology that encompasses two empirically relevant regimes: assistive AI that complements human selling and autonomous AI that substitutes for it. Our central equilibrium result is that the optimal commission rate is lower under both AI regimes than in the benchmark without AI. The mechanism is that AI weakens the connection between observed sales and the salesperson's contribution, making commissions less effective on both margins: each additional unit of commission elicits less salesperson-driven output, and performance pay attracts high-ability candidates less effectively because AI narrows the output gap between strong and weak salespeople. The practical implication is that firms should not simply layer AI onto legacy compensation plans; as AI's role grows, reliance on commission-based pay should be recalibrated.

Keywords: salesforce compensation; artificial intelligence; human–AI collaboration; adverse selection; moral hazard; contract theory

[1]The authors are listed alphabetically and contributed equally to this work. The authors thank helpful comments from Anthony Dukes, Zheng Gong, Jussi Keppo, Z. Eddie Ning; and from audiences at 2025 Informs Annual Conference, 2026 POMS-HK Conference, 2026 POMS Annual Conference, Singapore Rising Scholars Conference, and 2026 Marketing Science Conference.

[*]Associate Professor of Marketing, Hunan University, Business School, husaiquan@hnu.edu.cn.

[†]Ph.D. Candidate, National University of Singapore, Institute of Operations Research and Analytics, bo.pang@u.nus.edu.

[‡]Assistant Professor of Management, Shanghai University of International Business and Economics, School of Management, xugx@suibe.edu.cn.

[§]Assistant Professor of Marketing, National University of Singapore, NUS Business School, y.zhu@nus.edu.sg.

1 Introduction

Artificial intelligence (AI) is rapidly transforming how firms organize productive work. In McKinsey's most recent global survey of nearly 2,000 organizations across 105 countries, 88% of respondents report using AI in at least one business function (Singla et al. 2025). Official statistics corroborate this rapid diffusion: across OECD countries with available data, the share of firms reporting AI use more than doubled from 8.7% in 2023 to 20.2% in 2025, reaching 52.0% among large firms.[1] Newly available U.S. Census data show that 18% of firms had adopted AI by early 2026, and that these adopters account for 32% of U.S. employment (Bonney et al. 2026).

Yet the diffusion of AI has run far ahead of firms' understanding of how to organize around it. Challapally et al. (2025) finds that only about 5% of enterprise AI pilots generate measurable profit and loss impact. Strikingly, what separates the high performers is not superior technology but organizational adaptation (Singla et al. 2025). This pattern echoes a central insight from the economics of technology adoption: general purpose technologies create value only when accompanied by complementary organizational changes, including, but not limited to, changes in how work is structured, how responsibilities are allocated, and how workers are motivated. Such organizational investments, however, systematically lag behind technological progress (David 1990; Bresnahan and Trajtenberg 1995; Aral et al. 2012; Brynjolfsson et al. 2021). The key question, for each business function, is what the complementary organizational redesign should look like.

This paper takes up that question for the sales function, where AI adoption is already widespread (Bonney et al. 2026). In Salesforce's most recent survey of 4,050 sales professionals across 22 countries, 54% of sales teams already use AI agents deployed across the sales cycle, and nine in ten either use them or expect to within two years (Salesforce, Inc. 2026). As the core revenue-generation activity of the firm, selling has also emerged as a

[1] https://www.oecd.org/en/about/news/announcements/2026/01/ai-use-by-individuals-surges-across-the-oecd-as-adoption-by-firms-continues-to-expand.html

focal point of the AI-era labor market: three of the ten fastest-growing jobs in the United States are sales roles, according to LinkedIn's analysis of millions of job transitions between 2023 and 2025.[2]

Specifically, we study optimal contracting for the salesforce once AI enters the production of sales, along two margins. On the intensive margin, the contract must motivate a given salesperson to work harder; on the extensive margin, it determines what kind of salespeople the firm attracts. Understanding the second margin becomes essential in the presence of AI. By changing the technology through which effort and ability translate into sales, AI alters not only how hard a given salesperson works under a contract, but also which types of salespeople find that contract worth pursuing in the first place. These questions are pressing rather than prospective: salespeople are traditionally managed through explicit performance-based contracts, so the organization of selling work is tightly intertwined with the design of incentives. We deliver our insights through a game-theoretic model, and theory, we argue, is the right tool at this stage. The practitioner debate is vigorous but atheoretical, and managerial experimentation has outpaced systematic guidance. Meanwhile, credible empirical identification of the complete contracting problem remains difficult with currently available data.

We develop a parsimonious model that integrates both margins within a single framework. A firm seeks to fill a positive mass of sales positions by offering a linear compensation contract consisting of a base salary and a commission rate to a pool of candidates who privately know their own ability. The extensive margin operates through candidate self-selection: candidates differ in their outside opportunities, with high-ability candidates having systematically better outside opportunities, so the posted contract determines the ability mix of the applicant pool, while the firm's imperfect screening technology can only partially distinguish between candidate types. The intensive margin operates through post-hire effort: the hired salesperson, whose cost of selling effort depends on ability,

[2] https://www.linkedin.com/pulse/linkedin-jobs-rise-2026-25-fastest-growing-roles-us-linkedin-news-dlb1c

chooses how hard to work under the contract. Sales output is jointly produced by the salesperson and AI through a constant-elasticity-of-substitution (CES) production technology (Arrow et al. 1961). This flexible specification allows us to nest within a unified framework the two human–AI production regimes: one in which AI complements human selling effort and another in which AI substitutes for it.

In practice, AI enters sales production in two economically distinct forms. The first is assistive: AI serves as a tool that the salesperson wields, and its value is realized only through the salesperson's own effort. Prospecting platforms score leads and surface buying signals that tell a representative whom to contact and when; generative tools draft outreach messages and synthesize account research; and conversation intelligence systems monitor live sales calls and provide real-time guidance on objection handling and pricing (Salesforce, Inc. 2026). In each case, AI increases the productivity of human selling effort, but it produces nothing if the salesperson does not act: unread lead scores, unsent drafts, and ignored recommendations produce no revenue. Assistive AI is thus naturally modeled as a complement to human effort, because improvements in AI raise the marginal return to the salesperson's effort. The second form is autonomous: AI operates as an independent producer whose output arrives whether or not the salesperson exerts effort. Autonomous agents can contact, nurture, and qualify leads around the clock, including prospects that human sellers lack the capacity to pursue. Salesforce, for example, reports that its own agents contacted 130,000 previously untouched leads and created 3,200 sales opportunities within four months (Salesforce, Inc. 2026). In more fully automated settings, AI can conduct an entire sales interaction for relatively standardized products: identifying customer needs, personalizing an offer, generating a quote, and completing the transaction without a human salesperson. Because autonomous AI can generate positive sales output even when the salesperson exerts no effort, it is naturally modeled as a substitute for human effort, at least for the tasks or transactions it can perform independently.

Our central equilibrium result delivers a simple managerial message: whether AI complements or substitutes for human selling, firms should reduce the commission rate relative to the benchmark without AI. The common economic force is that AI weakens the link between observed sales and the salesperson's own contribution. A commission serves two organizational purposes. On the intensive margin, it motivates the hired salesperson to exert greater effort. On the extensive margin, it helps the firm attract high-ability candidates, who can convert strong incentives into sales more efficiently. AI weakens the effectiveness of commissions on both margins, though the specific channel differs across the two AI regimes, as we detail when presenting the model results. What is common to both regimes is that commissions become blunter on both margins. As an incentive instrument, each commission dollar buys less additional effort. As a selection instrument, it becomes less effective at distinguishing between candidate types. Although AI benefits both high- and low-ability salespeople, selection depends on their *relative* gains, and AI narrows the output gap between strong and weak salespeople. Thus, as AI assumes a larger role in sales production, commissions become a less targeted instrument for both motivating effort and attracting talent.

The practical implication is not that commissions cease to matter, or that firms should mechanically replace them with higher base salaries. Rather, firms should avoid layering AI onto legacy compensation plans without adjustment: as AI assumes a larger role in producing sales, each dollar of commission becomes a less targeted, and therefore more expensive, instrument for eliciting effort and attracting high-ability sellers.

The remainder of the paper proceeds as follows. We discuss how our paper relates to the existing literature in Section 2. Section 3 introduces the model, and Section 4 characterizes the equilibrium and develops the main comparative results across alternative human–AI production regimes. Section 5 discusses the scope and boundary conditions of our findings and concludes.

2 Literature Review

Our paper builds on the salesforce compensation literature initiated by Basu et al. (1985), who characterize the optimal compensation plan in an agency-theoretic framework. Subsequent research has extended this literature along several dimensions, including menus of contracts and nonlinear compensation (Lal and Staelin 1986; Rao 1990), quota-based schemes (Raju and Srinivasan 1996), monitoring and information acquisition (Joseph and Thevaranjan 1998), sales contests and relative-performance incentives (Kalra and Shi 2001), price delegation and channel contracting (Bhardwaj 2001; Kräkel and Schöttner 2020), dynamic incentive provision (Rubel and Prasad 2016), and double moral hazard (Dai and Jerath 2019). Building on the framework of Castro-Pires and Georgiadis (2026), we extend the salesforce compensation literature along two dimensions. First, we endogenize sales candidates' application decisions, thereby adding a pre-hire selection margin (Lo et al. 2011) to the standard post-hire incentive problem: compensation affects not only the effort exerted by the salesperson who is hired, but also the ability composition of the applicant pool and, ultimately, of the firm's hires. Second, we model sales as the output of a CES production technology combining human selling effort and AI input. This formulation nests economically distinct regimes of human–AI complementarity and substitutability, allowing us to characterize how the nature of human–AI interaction reshapes the optimal compensation plan.

Our paper also connects to a growing body of analytical work on the economics of AI. A foundational perspective conceptualizes AI as a technology that lowers the cost of prediction and emphasizes that automating prediction may either substitute for or complement human labor, depending on how prediction interacts with judgment and other decision tasks (Agrawal et al. 2019). Building on this perspective, one stream examines how AI reshapes the organization of human work. Athey et al. (2020) study the allocation of decision authority between humans and AI and show that a more capable AI can weaken human information-acquisition incentives, so a technically superior AI need not generate

greater organizational value. Dai and Singh (2020) examine how privately informed experts' use of diagnostic tools, including AI, interacts with reputational signaling, whereas Fügener et al. (2026) characterize when AI should automate tasks, augment human decision makers, or leave tasks to humans under different forms of human–AI complementarity. A parallel stream studies how AI and algorithms alter firms' product-market strategies and market outcomes. Using computational experiments, Calvano et al. (2020) show that autonomous pricing algorithms can learn collusive strategies; Cao et al. (2024) examine how competition shifts recommendation algorithms' exploration–exploitation tradeoff; Choi et al. (2024) show that advances in prediction capability can nonmonotonically alter a firm's pricing and surplus-extraction strategy; and Ning et al. (2025) demonstrate that more accurate prediction can induce strategic mistargeting and reshape consumers' privacy choices. Most closely related to our worker-heterogeneity mechanism, Zou et al. (2026) show that generative AI can disproportionately improve lower-quality creators, compressing quality differences and thereby changing equilibrium entry and welfare. We contribute to this literature by introducing AI into a principal–agent model of sales compensation in which the same contract governs both post-hire effort and pre-hire worker selection. This allows us to characterize how assistive and autonomous AI reshape the optimal commission rate and the endogenous composition of the sales force.

The empirical literature examining how AI reshapes business organizations remains nascent. Two empirical studies are particularly relevant to the mechanism underlying our analysis. Noy and Zhang (2023) study the productivity effects of ChatGPT in mid-level professional writing tasks, whereas Brynjolfsson et al. (2025) examine the introduction of a generative AI–based conversational assistant among customer-support agents. Despite the different task settings, both studies find that lower-skilled workers benefit more from AI than higher-skilled workers. This pattern is consistent with a central force in our model: by compressing the productivity gap between stronger and weaker sales candidates, AI reduces the extent to which performance-based pay distinguishes between

worker types and therefore weakens the effectiveness of commissions as a selection instrument.[3]

3 The Model

Consider a firm and a unit mass of potential sales candidates. Each candidate is characterized by a privately observed type $i \in \{H, L\}$, where H denotes a high-ability type and L denotes a low-ability type. Let $p \in (0, 1)$ denote the ex-ante proportion of high-ability candidates in the population.

The game begins with the firm posting a sufficiently small positive mass of sales positions to ensure that all positions are filled in equilibrium. The firm commits to a linear compensation scheme, $w(X) = \alpha + \beta X$, where X denotes realized sales output, $\alpha \geq 0$ is the fixed component (base salary), and $\beta \in [0, 1]$ captures the incentive intensity (commission rate).[4] Although sales organizations might also employ quotas and other nonlinear compensation plans, we focus on linear contracts both for tractability and because they are well-grounded in the literature. In canonical principal–agent models, linear contracts emerge as optimal (or approximately optimal) under standard assumptions (Basu et al. 1985; Holmstrom and Milgrom 1987). Moreover, empirical evidence in sales settings documents the widespread use and effectiveness of linear or locally linear incentive schemes (Chung et al. 2020).

Next, after observing the firm's announcement, each sales candidate decides whether to apply for the positions based on his type. Conditional on his type $i \in \{H, L\}$, a worker draws an outside option $\overline{u}_i$ from a type-specific distribution $G_i(\cdot)$. We assume that outside options are uniformly distributed. In particular, high-type candidates draw

[3]Sales work may differ in important ways from professional writing and customer support. Accordingly, whether AI similarly compresses productivity differences across salesperson types remains an empirical question.

[4]Empirical research shows that firms commonly offer uniform contracts even when their sales agents are heterogeneous (Daljord et al. 2016); we follow this empirical regularity.

$\overline{u}_i \sim$ Uniform$[0, U]$, where U is sufficiently large, whereas low-type candidates draw $\overline{u}_i \sim$ Uniform$[0, \eta U]$, with $\eta \in (0, 1)$. Under this specification, high-ability candidates possess weakly better outside opportunities than low-ability candidates. The outside-option distributions satisfy first-order stochastic dominance: $G_H(u) \leq G_L(u)$ for all u, with strict inequality for $u \in (0, \eta U)$. This assumption is consistent with empirical evidence suggesting that more qualified applicants typically enjoy stronger alternative employment opportunities (Cullen and Pakzad-Hurson 2023; Caldwell and Danieli 2024).

After receiving the applications, the firm conducts a screening process designed to assess candidates' sales aptitude. The screening technology is imperfect: a type-i candidate passes the screen with probability r_i. To capture systematic differences in sales capability, we assume $r_H > r_L$, implying that high-ability sales candidates are strictly more likely to pass the screening stage than low-ability candidates. Among the applicants who pass the screening stage, the firm randomly selects candidates to fill the positions. Random selection captures residual uncertainty and idiosyncratic considerations in hiring decisions, which is aligned with the labor literature (Lazear and Rosen 1981; Postel-Vinay and Robin 2002).

The hired sales representative then chooses an effort level $e_i \geq 0$, taking into account his type and the AI technology. Sales output, denoted by X_i, is jointly determined by the representative's selling effort and the AI tools. Effort is privately costly. A type-i salesperson incurs cost $\psi(e_i) = \frac{1}{2} c_i e_i^2$, where c_i captures the marginal cost of exerting selling effort. The quadratic specification ensures convex effort costs and is widely adopted in the sales compensation and agency literature (e.g., Rao 1990; Hauser et al. 1994; Joseph and Thevaranjan 1998). To capture heterogeneity in selling capability, we normalize $\frac{1}{2} < c_H = c < c_L = 1$, so that high-ability salespeople exert effort more efficiently than low-ability salespeople. The restriction $c \in \left(\frac{1}{2}, 1\right)$ rules out an extreme "superstar" type with arbitrarily low effort costs while preserving meaningful heterogeneity across sales representatives.

To model the production technology jointly determined by human effort and AI, we adopt a constant elasticity of substitution (CES) production function, which is widely used in macroeconomics (e.g., Arrow et al. 1961; Brown and Cani 1963; Klump and Preissler 2000) and has recently been adapted to analyze the role of data in production (Jones and Tonetti 2020; Chen and Keppo 2025). Specifically, we model the output function as

$$X_i = (\gamma\, x_i^{\rho} + (1-\gamma)\, \delta^{\rho})^{\frac{1}{\rho}} + \epsilon,$$

where $x_i \geq 0$ denotes the human input and $\delta \in (\underline{\delta}, \overline{\delta})$ represents the AI input, which is exogenously given. We restrict the AI input to an intermediate range so that AI is neither too weak to matter nor so productive that it renders human effort irrelevant. The additive normal specification follows the standard agency treatment of salesforce compensation (Holmstrom and Milgrom 1987; Lal and Srinivasan 1993; Joseph and Thevaranjan 1998). Under this specification, realized output may be negative; accordingly, X_i is interpreted as net sales contribution (gross revenue less returns, discounts, rebates, and selling costs) or as sales performance relative to a predetermined benchmark. For analytical tractability, we set $x_i = e_i \geq 0$ (Holmstrom and Milgrom 1991; Levitt 1995). The parameter $\gamma \in [0, 1]$ captures the degree of AI penetration in the production process. Lower values of γ correspond to greater AI involvement. We treat γ as exogenously given. The stochastic term $\epsilon \sim \mathcal{N}(0, \sigma^2)$ captures idiosyncratic outcome shocks or market-related uncertainty (Joseph and Thevaranjan 1998; Bolton and Dewatripont 2004; Mihm and Schlapp 2019; Long and Nasiry 2020).

A central parameter in the CES production function is $0 < \rho \leq 1$, which governs the interaction between a human sales representative and AI technology. ρ determines whether AI tools primarily amplify the effectiveness of human output or replace parts of the salesperson's role in generating revenue. Thus, it governs both the degree and the nature between human and AI inputs. Specifically:

- **As $\rho \to 0$ (Cobb–Douglas Form: Complementarity):** The CES function converges to the Cobb–Douglas form,

$$X_i = x_i^{\gamma} \delta^{1-\gamma} + \epsilon,$$

under which human effort and AI are complementary inputs: the cross-partial $\partial^2 X_i / \partial x_i \partial \delta = \gamma(1-\gamma) x_i^{\gamma-1} \delta^{-\gamma} \geq 0$, so better AI raises the marginal return to selling effort, and effort is essential. This regime captures assistive deployments.

- **When $\rho = 1$ (Linear Form: Substitutability):** The CES production function becomes linear:

$$X_i = \gamma x_i + (1-\gamma)\delta + \epsilon.$$

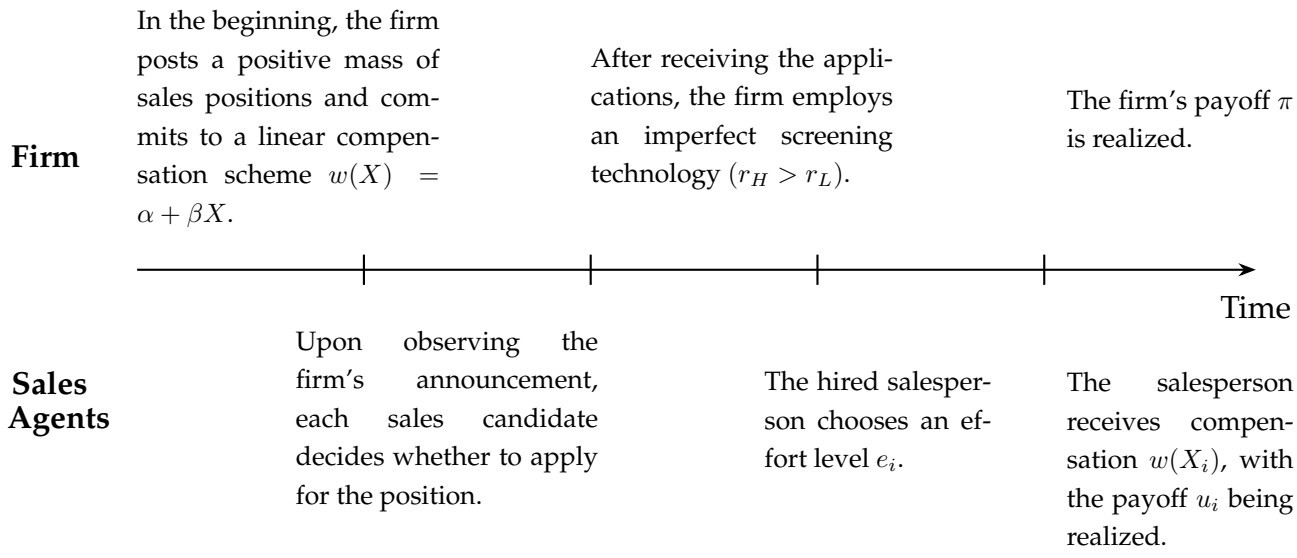
Here the cross-partial is zero, so AI leaves the marginal return to effort unchanged, and output no longer requires human effort. This regime captures autonomous deployments, in which AI operates as a parallel producer whose contribution flows into the salesperson's performance measures independently of his effort.[5]

After output X_i is realized, payoffs are determined. A hired salesperson receives compensation $w(X_i)$. We denote the firm's profit as $\pi(X_i, w) = X_i - w(X_i)$. We assume the salesperson is risk-neutral, with utility given by $u_i = w(X_i) - \psi(e_i)$ (e.g., Sappington 1983; Poblete and Spulber 2012; Simester and Zhang 2014). Sales representatives who are not hired receive the outside option $\overline{u}$. We summarize the sequence of our model in Figure 1 and list all notations in Table 1. We solve for Subgame Perfect Bayesian Equilibrium (SPBE) using backward induction[6].

[5] A natural question is why the firm compensates the salesperson on total output X_i rather than filtering out the AI component and rewarding $\gamma x_i + \epsilon$ alone. Filtering is redundant in our setting: because δ is deterministic, the contract (α, β) on total output is payoff-equivalent to a contract on filtered output with base salary $\alpha + \beta(1-\gamma)\delta$, since the AI component enters compensation only as a constant transfer that the firm undoes through the fixed payment. Our formulation is therefore without loss of generality with respect to attribution.

[6] We note that our model incorporates both moral hazard, through the unobserved effort of the hired salesperson, and adverse selection, through candidates' private information about their ability (e.g., Hu et al. 2016).

Figure 1. Sequence of Events



4 Equilibrium Analysis

We begin by characterizing the hired salespeople's effort choice and prospective salespeople's application decisions, illustrating how the firm's compensation scheme shapes the composition of its hires. We then characterize the optimal contract under three empirically relevant cases: no AI, AI complementarity, and AI substitutability. Finally, based on the equilibrium analysis, we compare these cases to examine how different AI regimes affect the firm's optimal contract, the skill composition of its workforce, and the expected output.

4.1 Salesperson's Problem

For ease of notation, define $f(e) \equiv (\gamma e^{\rho} + (1-\gamma)\delta^{\rho})^{1/\rho}$. The following lemma characterizes the hired salespeople's optimal effort choice.

Lemma 1. *Suppose that* $0 < \rho \leq 1$ *and* $\beta > 0$*. For a hired salesperson of type* $i \in \{H, L\}$*, the*

Table 1. Summary of Notations

Notation	Description
$i \in \{H, L\}$	Type of sales candidates: H for high-ability type, L for low-ability type
$p \in (0, 1)$	The ex-ante proportion of high-ability candidates in the population
$\overline{u}_i$	Outside option for type-i sales candidates
$G_i(\cdot)$	The cumulative distribution function of outside option for type-i sales candidates, $G_H \sim Uniform[0, U], G_L \sim Uniform[0, \eta U]$ with $\eta \in (0, 1)$
$r_i \in (0, 1)$	Probability that a type-i sales candidate passes the screening, $r_H > r_L$
X_i	The realized sales output of a type-i sales agent, where $X_i = (\gamma\, x_i^\rho + (1-\gamma)\, \delta^\rho)^{\frac{1}{\rho}} + \epsilon$
γ	The degree of AI penetration, $\gamma \in [0, 1]$
$x_i \geq 0$	Input measure of type-i sales agents, where $x_i = e_i \geq 0$
$\delta \in (\underline{\delta}, \overline{\delta})$	Input measure of AI system
ϵ	Outcome shock, where $\epsilon \sim \mathcal{N}(0, \sigma^2)$
$0 < \rho \leq 1$	Substitution parameter in the CES production function
$e_i \geq 0$	Effort level of a type-i sales agent
$\psi(e_i)$	The cost function of selling effort for type-i sales agents, where $\psi(e_i) = \frac{1}{2} c_i e_i^2$, c_i denotes the cost coefficient for type-i sales agents ($\frac{1}{2} < c_H = c < c_L = 1$)
$w(X_i)$	Compensation scheme as a function of output X_i, where $w(X_i) = \alpha + \beta X_i$ ($\alpha \geq 0$ denotes base salary, $\beta \geq 0$ denotes commission rate)
$u_i(w, e_i)$	Utility function of type-i sales agents: $u_i(w, e_i) = w(X_i) - \psi(e_i)$
$\pi(X_i, w)$	Firm's profit function: $\pi(X_i, w) = X_i - w(X_i)$

optimal effort level $e_i^* > 0$ is unique and is characterized by

$$\beta\gamma e_i^{*(\rho-2)} \left(\gamma e_i^{*\rho} + (1-\gamma)\delta^\rho\right)^{\frac{1-\rho}{\rho}} = c_i. \tag{1}$$

Proof. A hired salesperson of type $i \in \{H, L\}$ chooses effort to maximize expected utility:

$$\max_{e_i \geq 0} \{\mathbb{E}[w(X_i \mid e_i)] - \psi_i(e_i)\} = \max_{e_i \geq 0} \left\{ \alpha + \beta f(e_i) - \frac{1}{2} c_i e_i^2 \right\}.$$

For an interior solution, the first-order condition is

$$\beta\gamma e_i^{\rho-1} \left(\gamma e_i^{\rho} + (1-\gamma)\delta^{\rho}\right)^{\frac{1-\rho}{\rho}} - c_i e_i = 0.$$

Since the optimal effort level is strictly positive, dividing both sides by e_i yields Equation (1).

Moreover, the second derivative of the objective function is $\beta\gamma(\rho-1)(1-\gamma)\delta^{\rho} e_i^{\rho-2} f(e_i)^{1-2\rho} - c_i < 0$, where the inequality follows from $0 < \rho \leq 1$ and $c_i > 0$. Therefore, the salesperson's objective function is strictly concave in effort, and the first-order condition characterizes the unique global optimum. □

The optimal effort level e_i^* exhibits the standard comparative-static properties commonly found in the sales-force compensation literature. First, effort increases with the commission rate β. Totally differentiating the first-order condition with respect to β yields the results. That is, a higher commission rate increases the salesperson's marginal return to effort and induces greater effort provision. Second, optimal effort decreases with the salesperson's effort-cost parameter $e_H^* > e_L^*$. This can be derived by differentiating the first-order condition with respect to c_i. Finally, since $f(e)$ is strictly increasing in effort, we also have $f_H \equiv f(e_H^*) > f(e_L^*) \equiv f_L$. Hence, the high-type salesperson not only exerts more effort but also generates a higher expected sales outcome.

The next lemma characterizes candidates' application decisions. Let $u_i^* \equiv \mathbb{E}u_i(e_i^*)$ denote the expected utility obtained by a type-i candidate if hired.

Lemma 2. *Suppose that U is sufficiently large, then each sales candidate follows a cutoff rule for application. Specifically, a type-i candidate applies if and only if $\overline{u}_i \leq u_i^*$. The type-specific*

application probabilities are

$$q_H = \frac{u_H^*}{U}, \qquad q_L = \frac{u_L^*}{\eta U}. \tag{2}$$

Proof. Because $c_H = c \in (1/2, 1)$, $c_L = 1$ and U is sufficiently large, we can show that the associated utilities u_H^* and u_L^* are well defined and finite.

Consider a type-i candidate with outside option $\overline{u}_i$. If the candidate does not apply, he receives $\overline{u}_i$ with certainty. If he applies, let $h_i > 0$ denote his probability of being hired. The probability h_i may depend on the screening technology and on the application decisions of other candidates. If hired, the candidate receives utility u_i^*; if not hired, he retains his outside option $\overline{u}_i$. His expected utility from applying is therefore $h_i u_i^* + (1 - h_i)\overline{u}_i$. By contrast, his utility from not applying is $\overline{u}_i$. The net gain from applying is thus $h_i\big(u_i^* - \overline{u}_i\big)$. Since $h_i > 0$, a type-i candidate prefers to apply if and only if $\overline{u}_i \leq u_i^*$. Given $\overline{u}_H \sim \text{Uniform}[0, U]$ and $\overline{u}_L \sim \text{Uniform}[0, \eta U]$, the application probabilities are $q_H = u_H^*/U$ and $q_L = u_L^*/(\eta U)$, which proves Equation (2). □

Based on Lemma 2, we examine how the compensation scheme affects the ability composition of the firm's hires. Lemma 3 summarizes the results.

Lemma 3. *Suppose that U is sufficiently large and that $\beta > 0$. Then:*

1. *The probability that the hired salesperson is of type H is*

$$\mu \;=\; \frac{p\, q_H\, r_H}{p\, q_H\, r_H + (1-p)\, q_L\, r_L}. \tag{3}$$

2. *The effects of the compensation parameters on the ability composition of the hired salesforce are as follows:*

$$\frac{\partial \mu}{\partial \alpha} < 0, \qquad \textit{and} \qquad \frac{\partial \mu}{\partial \beta} > 0 \quad \Longleftrightarrow \quad \frac{f_H}{f_L} > \frac{q_H}{\eta q_L}. \tag{4}$$

Proof. Because the firm selects uniformly at random from among the candidates who pass, the posterior probability that a hired salesperson is of type H equals the proportion of type-H candidates among the effective applicants, as given by Equation (3).

It follows directly from Equation (3) that μ is strictly increasing in the ratio q_H/q_L. Moreover, Lemma 2 implies that $q_H/q_L = \eta u_H^*/u_L^*$. Therefore, to determine how the base salary and commission rate affect the ability composition of the hired salesforce, it suffices to examine how the ratio u_H^*/u_L^* varies with the compensation parameters.

By the envelope theorem, we have $\partial u_i^*/\partial\alpha = 1$ and $\partial u_i^*/\partial\beta = f_i^*$. We next establish that $u_H^* > u_L^*$. A type-H salesperson can choose the type-L salesperson's optimal effort while incurring a strictly lower effort cost:

$$u_H^* \geq \alpha + \beta f\big(e_L^*\big) - \frac{1}{2}\, c_H\big(e_L^*\big)^2 > \alpha + \beta f\big(e_L^*\big) - \tfrac{1}{2}\, c_L\big(e_L^*\big)^2 = u_L^*.$$

The first inequality follows because e_L^* is feasible for type H, whereas the second is strict because $c_H < c_L$ and $e_L^* > 0$.

Consequently,

$$\frac{\partial}{\partial\alpha}\left(\frac{u_H^*}{u_L^*}\right) = \frac{u_L^* - u_H^*}{(u_L^*)^2} < 0, \qquad \frac{\partial}{\partial\beta}\left(\frac{u_H^*}{u_L^*}\right) = \frac{f_H u_L^* - f_L u_H^*}{(u_L^*)^2}.$$

Thus, together with the monotonicity of μ in q_H/q_L, we have $\partial\mu/\partial\alpha < 0$ and $\partial\mu/\partial\beta > 0$ if and only if $f_H/f_L > u_H^*/u_L^* = q_H/(\eta q_L)$. □

Lemma 3 highlights that compensation shapes not only the hired salespeople's effort but also the composition of the workforce, that is, *who gets hired*. Moreover, the two compensation instruments affect this selection margin asymmetrically.

Because salesforce composition depends on the ratio of the two types' application rates, what matters for selection is not how much a compensation change raises each type's equilibrium utility in *absolute* terms, but how much it raises that utility *relative* to what the type already earns. This observation immediately delivers the first result. A higher base salary α raises both types' utilities by the same absolute amount; since $u_L^* < u_H^*$, this common increment constitutes a larger relative gain for low types, who therefore respond more strongly at the application margin.

The effect of the commission rate on salesforce composition is more subtle. Since $\partial u_i^*/\partial\beta = f_i$ and $f_H > f_L$, a higher commission confers a larger absolute benefit on high types. However, this comparison alone does not determine the direction of selection. Accordingly, the condition in Lemma 3 compares the relative marginal benefit of the commission, f_H/f_L, with the adjusted relative application rate, $q_H/(\eta q_L)$. When high types already dominate the applicant pool and their marginal earnings advantage is modest, a higher commission rate can *reduce* the probability that the hired salesperson is of high ability: the additional commission dollars are largely inframarginal for high types, while still drawing in new low types at the margin.

4.2 Optimal Sales Incentives

We now turn to the firm's optimal contract. To economize on notation, the preceding lemmas suppressed the dependence of equilibrium objects on the contract; we now make this dependence explicit, since the firm's problem is precisely to choose (α, β) anticipating how these objects respond.

The firm's expected profit has two components: the probability that the position is filled and the margin earned conditional on filling it. The first component is immediate based on our model assumptions. The firm's problem therefore reduces to maximizing the margin per hire.

For the second component, note that the induced effort $e_i^*(\beta)$, and hence expected output $f_i(\beta)$, depends on the contract only through the commission rate, whereas the pool composition $\mu(\alpha, \beta)$ of Lemma 3 depends on both instruments. Conditional on a hire, the firm pays the wage $\alpha + \beta X_i$ and collects X_i, so its expected margin per hire is

$$m(\alpha,\beta) \;=\; (1-\beta)\,\bar{f}(\alpha,\beta) - \alpha, \qquad \bar{f}(\alpha,\beta) \;\equiv\; \mu(\alpha,\beta)\,f_H(\beta) + \big(1-\mu(\alpha,\beta)\big)\,f_L(\beta),$$

where $\bar{f}$ is the expected output averaged over the equilibrium composition of the hired

pool. The contract (α, β) thus operates on profit through two distinct channels: the *selection* channel, via the pool composition μ entering the average output $\bar{f}$; and the *incentive* channel, via the induced efforts entering $f_i(\beta)$.

Deriving closed-form solutions over the general parameter space is difficult. We therefore focus on three empirically relevant cases to examine how AI shapes incentive structures through the two channels: no AI, AI that complements salespeople's productivity ($\rho \to 0$), and AI that substitutes for salespeople ($\rho = 1$). We characterize the optimal contracts in Proposition 1 and provide the proof in the Appendix.

Proposition 1. *Suppose that AI is neither too weak to matter nor so productive that it renders human effort irrelevant, $\delta \in (\underline{\delta}, \overline{\delta})$. The optimal contract is characterized as follows.*

1. ***No AI ($\gamma = 1$).*** *The optimal base salary is $\alpha_N^* = 0$ and the optimal commission rate is $\beta_N^* = 1/2$.*

2. ***AI complementarity ($\rho \to 0$).*** *The optimal base salary is $\alpha_C^* = 0$ and the optimal commission rate is $\beta_C^* = \gamma/2$.*

3. ***AI substitutability ($\rho = 1$).*** *The optimal base salary is $\alpha_S^* = 0$ and the optimal commission rate is the $\beta_S^* \in (0, 1)$ that satisfies*

$$\big(1-\beta_S^*\big)\left[\gamma^2\Big(\frac{\mu}{c}+1-\mu\Big) \;+\; \mu(1-\mu)\,\big(f_H-f_L\big)\Big(\frac{f_H}{u_H^*}-\frac{f_L}{u_L^*}\Big)\right] \;=\; \mu f_H + (1-\mu)\,f_L,$$

where all equilibrium objects are evaluated at $(\alpha, \beta) = (0, \beta_S^*)$, with $f_i = \gamma^2\beta/c_i + (1-\gamma)\delta$, $u_i^* = \beta\big(\gamma^2\beta/(2c_i) + (1-\gamma)\delta\big)$, and $\mu = \Big[\,1 + ((1-p)\,r_L)/(\eta\,p\,r_H)\,u_L^*/u_H^*\,\Big]^{-1}$.

4.3 Role of AI

Thus far, we have characterized the equilibrium under different AI regimes. In this section, we build on these equilibrium characterizations to compare the three cases and examine how different AI regimes affect the firm's optimal contract, workforce skill com-

position, and the expected output. Corollary 1 summarizes the results and the proofs are provided in the Appendix.

Corollary 1. *Compared with the no-AI benchmark, at the respective optimal contracts:*

1. *AI does not change the optimal base salary, which remains at zero, but it weakens incentive pay: both forms of AI feature strictly lower-powered commission rates than the no-AI benchmark.*

2. *Both forms of AI strictly dilute the quality of the hired pool: the probability that a hired salesperson is of high ability is highest in the no-AI benchmark.*

3. *AI strictly compresses the expected output gap between high- and low-ability salespeople.*

We now provide the economic intuition behind Corollary 1. First, neither AI regime changes the optimal base salary. Because α does not affect the hired salespeople's effort choice, its only role is selectional. As Lemma 3 shows, a higher base salary tilts the applicant pool toward low types. So under assumption $\alpha^* = 0$ in all regimes.

Second, consider the commission rate. A commission operates through two margins, incentives and selection, and both depend on output reflecting the salesperson's contribution: effort on the incentive margin and ability on the selection margin. AI erodes this premise under either regime, though through different channels. When AI substitutes for human effort, part of every sale is generated independently of the salesperson, so realized output partially reflects the machine's contribution rather than the individual's. When AI complements human effort, effort remains productive, but diminishing returns imply that differences in effort translate into less-than-proportional differences in sales. In both cases, output becomes a noisier signal of who the salesperson is and what they did. This blunts the commission on both margins. As an incentive device, each commission dollar buys less additional effort, since payments scale with output while the salesperson is responsible for only part of it. As a selection device, the commission loses discriminating power: when AI contributes to sales regardless of who makes them, the gap that

output-based pay creates between strong and weak salespeople narrows. With both the motivating and the sorting value of output-based pay reduced, the firm leans less on the commission.

Lastly, the output-gap and composition results follow from the same mechanism. Because AI contributes to sales regardless of who makes them, realized sales become less reflective of the salesperson's own ability, and the high type's relative output advantage f_H/f_L shrinks. Since pay is tied to output, this compression passes through to rewards, raising the equilibrium utility ratio u_L^*/u_H^*. By Lemmas 2 and 3, the probability of hiring a high type is strictly decreasing in this ratio, so composition deteriorates.

5 Discussions and Conclusion

This paper studies how salesforce compensation should adapt when AI enters the production of sales. In a model where a posted linear contract governs both how hard the hired salesperson works and which candidates apply, we show that whether AI assists human effort or produces output autonomously, the firm optimally reduces the commission rate relative to the no-AI benchmark. The common force is that AI weakens the connection between observed sales and the salesperson's own contribution, so output-based pay buys less effort and loses discriminating power in attracting talent, and the expected ability of the hired salesforce declines as a consequence. We close by discussing the scope of these results and the research they invite.

From the firm's perspective, our composition result is a pessimistic one: because AI narrows the output gap between high- and low-ability salespeople, a pattern consistent with emerging empirical evidence, the ability composition of the salesforce deteriorates in equilibrium. Two limitations qualify this conclusion and point to future work. First, we take the competitive environment as given through the outside-option distribution. This allows us to isolate the role of AI in reshaping the compensation plan of a single firm,

but it abstracts from how competition among firms would determine outside options in equilibrium; how AI adoption changes the structure of compensation across competing firms is an interesting direction for future research. Second, our model features no learning: each candidate's type is fixed. Yet AI can also act as a coach (Luo et al. 2021) that may transform low-ability salespeople into high-ability ones, and recent theory shows that AI reshapes how junior workers learn from senior workers (Ide and Talamàs 2025). Whether, when, and for whom AI upgrades ability rather than merely augmenting output is, in our view, a first-order question for empirical research, and its answer would determine how much of the composition deterioration we identify is offset over time.

References

Agrawal, A., J. S. Gans, and A. Goldfarb (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives 33*(2), 31–50.

Aral, S., E. Brynjolfsson, and L. Wu (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science 58*(5), 913–931.

Arrow, K. J., H. B. Chenery, B. Minhas, and R. M. Solow (1961). Capital-labor substitution and economic efficiency. *The Review of Economics and Statistics 43*(3), 225–250.

Athey, S. C., K. A. Bryan, and J. S. Gans (2020). The allocation of decision authority to human and artificial intelligence. *AEA Papers and Proceedings 110*, 80–84.

Basu, A. K., R. Lal, V. Srinivasan, and R. Staelin (1985). Salesforce compensation plans: An agency theoretic perspective. *Marketing Science 4*(4), 267–291.

Bhardwaj, P. (2001). Delegating pricing decisions. *Marketing science 20*(2), 143–169.

Bolton, P. and M. Dewatripont (2004). *Contract theory*. MIT press.

Bonney, K., C. Breaux, E. Dinlersoz, L. Foster, J. Haltiwanger, and A. Pande (2026). The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks. CES Working Paper 26-25, U.S. Census Bureau, Center for Economic Studies.

Bresnahan, T. F. and M. Trajtenberg (1995). General purpose technologies: 'engines of growth'? *Journal of Econometrics 65*(1), 83–108.

Brown, M. and J. S. D. Cani (1963). Technological change and the distribution of income. *International Economic Review 4*(3), 289–309.

Brynjolfsson, E., D. Li, and L. Raymond (2025, 05). Generative ai at work. *The Quarterly Journal of Economics 140*(2), 889–942.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics 13*(1), 333–372.

Caldwell, S. and O. Danieli (2024). Outside options in the labour market. *Review of Economic Studies 91*(6), 3286–3315.

Calvano, E., G. Calzolari, V. Denicolò, and S. Pastorello (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review 110*(10), 3267–3297.

Cao, H. H., L. Ma, Z. E. Ning, and B. Sun (2024). How does competition affect exploration vs. exploitation? a tale of two recommendation algorithms. *Management Science 70*(2), 1029–1051.

Castro-Pires, H. and G. Georgiadis (2026). Incentives and selection.

Challapally, A., C. Pease, R. Raskar, and P. Chari (2025). The GenAI divide: State of AI in business 2025. Technical report, MIT NANDA, Massachusetts Institute of Technology.

Chen, Z. and J. Keppo (2025). R&D data sharing in new product development. *Manufacturing & Service Operations Management 27*(4), 1275–1294.

Choi, W. J., Q. Liu, and J. Shin (2024). Predictive analytics and ship-then-shop subscription. *Management Science 70*(2), 1012–1028.

Chung, D. J., B. Kim, and N. B. Syam (2020). A practical approach to sales compensation: What do we know now? What should we know in the future? *Foundations and Trends in Marketing 14*(1), 1–52.

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

Dai, T. and K. Jerath (2019). Salesforce contracting under uncertain demand and supply: Double moral hazard and optimality of smooth contracts. *Marketing Science 38*(5), 852–870.

Dai, T. and S. Singh (2020). Conspicuous by its absence: Diagnostic expert testing under uncertainty. *Marketing Science 39*(3), 540–563.

Daljord, O., S. Misra, and H. S. Nair (2016). Homogeneous contracts for heterogeneous agents: Aligning sales force composition and compensation. *Journal of Marketing Research 53*(2), 161–182.

David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review 80*(2), 355–361.

Fügener, A., D. D. Walzner, and A. Gupta (2026). Roles of artificial intelligence in collaboration with humans: Automation, augmentation, and the future of work. *Management Science 72*(1), 538–557.

Hauser, J. R., D. I. Simester, and B. Wernerfelt (1994). Customer satisfaction incentives. *Marketing science 13*(4), 327–350.

Holmstrom, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55*(2), 303–328.

Holmstrom, B. and P. Milgrom (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization 7*, 24–52.

Hu, Y. J., J. Shin, and Z. Tang (2016). Incentive problems in performance-based online advertising pricing: Cost per click vs. cost per action. *Management Science 62*(7), 2022–2038.

Ide, E. and E. Talamàs (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy 133*(12), 3762–3800.

Jones, C. I. and C. Tonetti (2020). Nonrivalry and the economics of data. *American Economic Review 110*(9), 2819–2858.

Joseph, K. and A. Thevaranjan (1998). Monitoring and incentives in sales organizations: An agency-theoretic perspective. *Marketing Science 17*(2), 107–123.

Kalra, A. and M. Shi (2001). Designing optimal sales contests: A theoretical perspective. *Marketing Science 20*(2), 170–193.

Klump, R. and H. Preissler (2000). CES production functions and economic growth. *The Scandinavian Journal of Economics 102*(1), 41–56.

Kräkel, M. and A. Schöttner (2020). Delegating pricing authority to sales agents: The impact of kickbacks. *Management Science 66*(6), 2686–2705.

Lal, R. and V. Srinivasan (1993). Compensation plans for single-and multi-product salesforces: An application of the holmstrom-milgrom model. *Management Science 39*(7), 777–793.

Lal, R. and R. Staelin (1986). Salesforce compensation plans in environments with asymmetric information. *Marketing science 5*(3), 179–198.

Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–864.

Levitt, S. D. (1995). Optimal incentive schemes when only the agents'" best" output matters to the principal. *The RAND Journal of Economics 26*(4), 744–760.

Lo, D., M. Ghosh, and F. Lafontaine (2011). The incentive and selection roles of sales force compensation contracts. *Journal of Marketing Research 48*(4), 781–798.

Long, X. and J. Nasiry (2020). Wage transparency and social comparison in sales force compensation. *Management Science 66*(11), 5290–5315.

Luo, X., M. S. Qin, Z. Fang, and Z. Qu (2021). Artificial intelligence coaches for sales agents: Caveats and solutions. *Journal of Marketing 85*(2), 14–32.

Mihm, J. and J. Schlapp (2019). Sourcing innovation: On feedback in contests. *Management science 65*(2), 559–576.

Ning, Z. E., J. Shin, and J. Yu (2025). Targeted advertising as implicit recommendation: Strategic mistargeting and personal data opt-out. *Marketing Science 44*(2), 390–410.

Noy, S. and W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science 381*(6654), 187–192.

Poblete, J. and D. Spulber (2012). The form of incentive contracts: agency with moral hazard, risk neutrality, and limited liability. *The RAND Journal of Economics 43*(2), 215–234.

Postel-Vinay, F. and J.-M. Robin (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica 70*(6), 2295–2350.

Raju, J. S. and V. Srinivasan (1996). Quota-based compensation plans for multiterritory heterogeneous salesforces. *Management Science 42*(10), 1454–1462.

Rao, R. C. (1990). Compensating heterogeneous salesforces: Some explicit solutions. *Marketing Science 9*(4), 319–341.

Rubel, O. and A. Prasad (2016). Dynamic incentives in sales force compensation. *Marketing Science 35*(4), 676–689.

Salesforce, Inc. (2026). State of sales, seventh edition. Technical report, Salesforce Research.

Sappington, D. (1983). Limited liability contracts between principal and agent. *Journal of economic Theory 29*(1), 1–21.

Simester, D. and J. Zhang (2014). Why do salespeople spend so much time lobbying for low prices? *Marketing Science 33*(6), 796–808.

Singla, A., A. Sukharevsky, B. Hall, L. Yee, M. Chui, and T. Balakrishnan (2025). The state of AI in 2025: Agents, innovation, and transformation. Technical report, McKinsey & Company, QuantumBlack.

Zou, T., Z. Shi, and Y. Wu (2026). Welfare implications of democratization in content creation: Generative AI and beyond. *Journal of Marketing Research 63*(4), 607–625.

Appendix

Proof of Proposition 1:

We first show that the optimal base salary is zero in all cases. The firm chooses (α, β) to maximize its margin $m(\alpha, \beta) = (1-\beta)\big[\mu f_H + (1-\mu) f_L\big] - \alpha$. Because effort and hence f_i is independent of α,

$$\frac{\partial m}{\partial \alpha} = (1-\beta)\big(f_H - f_L\big)\frac{\partial \mu}{\partial \alpha} - 1 < 0,$$

where the inequality follows from $f_H > f_L$ and Lemma 3, which establishes that $\partial\mu/\partial\alpha < 0$. Therefore, every optimal contract sets the base salary at its lower bound, $\alpha^* = 0$. It remains to characterize the optimal commission rate conditional on $\alpha^* = 0$.

We begin with the no-AI case. Setting $\gamma = 1$, Lemma 1 implies that $e_i^* = \beta/c_i$. Hence, $f_i = \beta/c_i$ and $u_i^* = \beta^2/(2c_i)$. By Lemmas 2 and 3, the probability that a hired salesperson is of type H is

$$\mu = \left[1 + \frac{(1-p)r_L}{\eta p r_H}\frac{u_L^*}{u_H^*}\right]^{-1} = \left[1 + \frac{(1-p)r_L}{\eta p r_H}\frac{c_H}{c_L}\right]^{-1} \equiv \mu_N.$$

Thus, the composition is constant and independent of the contract. The firm's problem can therefore be written as

$$\max_{\beta}\ m(0,\beta) = (1-\beta)\left(\mu_N \frac{\beta}{c} + (1-\mu_N)\beta\right).$$

Recall that $\frac{1}{2} < c_H = c < c_L = 1$. Because the expression in parentheses is linear in β, the firm's objective is proportional to $\beta(1-\beta)$. It follows that

$$\beta_N^* = \frac{1}{2}.$$

Next, consider the AI-complementarity case. As $\rho \to 0$, the CES production technology converges to the Cobb–Douglas form $f(e,\delta) = e^{\gamma}\delta^{1-\gamma}$. The optimal effort of

a hired type-i salesperson is therefore $e_i^*(\beta) = (\beta\gamma\delta^{1-\gamma}/c_i)^{\frac{1}{2-\gamma}}$. The expected output is $f_i(\beta) = (\beta\gamma/c_i)^{\frac{\gamma}{2-\gamma}}\,\delta^{\frac{2(1-\gamma)}{2-\gamma}}$ and the utility is $u_i^*(\beta) = ((2-\gamma)/2)\beta f_i$.

As before, the probability that a hired salesperson is of type H is

$$\mu = \left[1 + \frac{(1-p)r_L}{\eta p r_H}\frac{u_L^*}{u_H^*}\right]^{-1} = \left[1 + \frac{(1-p)r_L}{\eta p r_H}c^{\frac{\gamma}{2-\gamma}}\right]^{-1} \equiv \mu_C.$$

Thus, as in the no-AI case, the composition of the pool is independent of the contract. The firm's problem can therefore be written as

$$m(0,\beta) = \left[\mu_C c^{-\frac{\gamma}{2-\gamma}} + 1 - \mu_C\right]\gamma^{\frac{\gamma}{2-\gamma}}\delta^{\frac{2(1-\gamma)}{2-\gamma}}\beta^{\frac{\gamma}{2-\gamma}}(1-\beta).$$

Maximizing this expression with respect to β yields

$$\beta_C^* = \frac{\gamma}{2}.$$

Finally, consider the AI-substitution case. When $\rho = 1$, the production function is linear $f(e,\delta) = \gamma e + (1-\gamma)\delta$. It follows that $e_i^*(\beta) = \beta\gamma/c_i$, $f_i(\beta) = \gamma^2\beta/c_i + (1-\gamma)\delta$, and $u_i^*(\beta) = \beta\left(\gamma^2\beta/(2c_i) + (1-\gamma)\delta\right)$.

In this case, the probability that a hired salesperson is of type H is no longer invariant to the sales contract. To characterize the optimal commission rate, we first derive

$$\frac{\partial\mu}{\partial\beta} = \mu(1-\mu)\left[\frac{\partial u_H^*/\partial\beta}{u_H^*} - \frac{\partial u_L^*/\partial\beta}{u_L^*}\right] = \mu(1-\mu)\left(\frac{f_H}{u_H^*} - \frac{f_L}{u_L^*}\right), \tag{A.1}$$

where the second equality follows from the envelope theorem.

Differentiating the firm's objective with respect to the commission rate yields

$$\frac{\partial m(0,\beta)}{\partial\beta} = -\bar{f}(\beta) + (1-\beta)\left[\mu\frac{\gamma^2}{c} + (1-\mu)\gamma^2 + (f_H - f_L)\frac{\partial\mu}{\partial\beta}\right].$$

Substituting (A.1), the stationarity condition $\partial m(0,\beta)/\partial\beta = 0$ becomes

$$(1-\beta)\left[\gamma^2\left(\frac{\mu}{c}+1-\mu\right)+\mu(1-\mu)(f_H-f_L)\left(\frac{f_H}{u_H^*}-\frac{f_L}{u_L^*}\right)\right]=\mu f_H+(1-\mu)f_L,$$

which is precisely condition (3) in the proposition.

As $\beta \to 0^+$, the composition of the pool converges to $\mu_0 \equiv (\eta p r_H)/(\eta p r_H + (1-p)r_L)$. Therefore,

$$\lim_{\beta\to 0^+}\frac{\partial m(0,\beta)}{\partial\beta}=\gamma^2\left(\frac{\mu_0}{c}+1-\mu_0\right)-(1-\gamma)\delta>0,$$

where the inequality follows from the assumption that $\delta < \overline{\delta}$. Hence, there exists some $\beta' \in (0,1)$ such that $m(0,\beta') > (1-\gamma)\delta \geq m(0,0^+)$, so the optimum cannot lie at the lower boundary. Nor can it occur at $\beta = 1$, since $m(0,1) = 0$. By continuity, a maximizer β_S^* exists and lies in the interior of $(0,1)$. It therefore satisfies the first-order condition. Moreover, feasibility is ensured because $u_i^*(\beta_S^*) > 0$.

Proof of Corollary 1:

Proposition 1 establishes that

$$\alpha_N^* = \alpha_C^* = \alpha_S^* = 0, \qquad \beta_N^* = \frac{1}{2}, \qquad \beta_C^* = \frac{\gamma}{2} < \frac{1}{2},$$

where the final inequality follows from $\gamma < 1$. It remains to prove that $\beta_S^* < 1/2$.

The key step is to show that, in the AI-substitute case,

$$\frac{\partial m(0,\beta)}{\partial\beta} < 0 \qquad \text{for every } \beta \in [1/2, 1].$$

It then follows that the maximizer must lie strictly below $1/2$. To establish this result, we bound separately the incentive and selection channels. Following Proposition 1 and its

proof, define the incentive channel by

$$I(\beta) \equiv \mu f_H'(\beta) + (1-\mu) f_L'(\beta) = \gamma^2 \left(\frac{\mu}{c} + 1 - \mu \right),$$

and the selection channel by

$$S(\beta) \equiv (f_H - f_L)\mu'(\beta) = \mu(1-\mu)(f_H - f_L) \left(\frac{f_H}{u_H^*} - \frac{f_L}{u_L^*} \right).$$

The proof of Proposition 1 gives

$$\frac{\partial m(0,\beta)}{\partial \beta} = -\bar{f}(\beta) + (1-\beta)\big[I(\beta) + S(\beta)\big], \qquad \bar{f}(\beta) = \beta I(\beta) + (1-\gamma)\delta,$$

where the expression for $\bar{f}$ follows from the expression of f_i. For $\beta \geq 1/2$, we have $(1-\beta)I(\beta) \leq I(\beta)/2 \leq \beta I(\beta)$. Consequently,

$$\frac{\partial m(0,\beta)}{\partial \beta} \leq (1-\beta)S(\beta) - (1-\gamma)\delta \leq \frac{S(\beta)}{2} - (1-\gamma)\delta.$$

It remains to bound the selection channel. From the proof of Proposition 1, $f_H - f_L = \gamma^2 \beta \left(1/c - 1\right)$ and $f_H u_L^* - f_L u_H^* = \gamma^2 \beta^2 (1-\gamma)\delta \left(1/c - 1\right)/2$. Moreover, $\mu(1-\mu) \leq 1/4$, and

$$u_H^* u_L^* = \beta^2 \left(\frac{\gamma^2 \beta}{2c} + (1-\gamma)\delta \right) \left(\frac{\gamma^2 \beta}{2} + (1-\gamma)\delta \right) > \beta^2 \cdot \frac{\gamma^4 \beta^2}{4c}.$$

Therefore,

$$S(\beta) = \mu(1-\mu)\frac{(f_H - f_L)(f_H u_L^* - f_L u_H^*)}{u_H^* u_L^*} < \frac{1}{4}\frac{\gamma^4 \beta^3 \left(\frac{1}{c} - 1\right)^2 (1-\gamma)\delta/2}{\gamma^4 \beta^4/(4c)} = \frac{(1-c)^2}{2c\beta}(1-\gamma)\delta.$$

Because $\beta \geq 1/2$ and $c > 1/2$,

$$S(\beta) < \frac{(1-c)^2}{c}(1-\gamma)\delta < \frac{1}{2}(1-\gamma)\delta.$$

It follows that, for every $\beta \in [1/2, 1]$,

$$\frac{\partial m(0,\beta)}{\partial \beta} \leq \frac{S(\beta)}{2} - (1-\gamma)\delta < 0.$$

Thus $m(0,\beta)$ is strictly decreasing on $[1/2, 1]$, so

$$\beta_S^* < \frac{1}{2} = \beta_N^*.$$

We next establish the composition results. By Lemmas 2 and 3, the probability that a hired salesperson is of type H is

$$\mu = \left[1 + \frac{(1-p)r_L}{\eta p r_H}\frac{u_L^*}{u_H^*}\right]^{-1}.$$

Because μ is strictly decreasing in u_L^*/u_H^*, it suffices to show that this ratio is strictly smaller in the no-AI benchmark than in either AI case.

In the no-AI benchmark, the proof of Proposition 1 gives $u_i^* = \beta^2/(2c_i)$. Thus,

$$\left(\frac{u_L^*}{u_H^*}\right)_N = \frac{c_H}{c_L} = c.$$

In the AI-complement case, $u_i^* = ((2-\gamma)/2)\,\beta f_i$ with $f_i \propto c_i^{-\gamma/(2-\gamma)}$. It follows that

$$\left(\frac{u_L^*}{u_H^*}\right)_C = c^{\gamma/(2-\gamma)} > c = \left(\frac{u_L^*}{u_H^*}\right)_N,$$

where the strict inequality follows because $c \in (1/2, 1)$ and $\gamma/(2-\gamma) \in (0, 1)$.

In the AI-substitute case, for any $\beta > 0$,

$$\left(\frac{u_L^*}{u_H^*}\right)_S - c = \frac{\gamma^2\beta/2 + (1-\gamma)\delta}{\gamma^2\beta/(2c) + (1-\gamma)\delta} - c = \frac{(1-c)(1-\gamma)\delta}{\gamma^2\beta/(2c) + (1-\gamma)\delta} > 0.$$

Therefore, we conclude that

$$\mu_C^* < \mu_N^* \qquad \text{and} \qquad \mu_S^* < \mu_N^*.$$

Finally, we establish the compression of the output gap. We compare the output ratio f_H/f_L across the three cases, each evaluated at its optimal contract.

In the no-AI benchmark, the proof of Proposition 1 gives $f_i = \beta/c_i$, so at any commission rate, and in particular at β_N^*,

$$\left(\frac{f_H}{f_L}\right)_N = \frac{c_L}{c_H} = \frac{1}{c}.$$

The equilibrium output gap without AI thus equals the full productivity gap between the types.

In the AI-complement case, $f_i \propto c_i^{-\gamma/(2-\gamma)}$, so

$$\left(\frac{f_H}{f_L}\right)_C = c^{-\frac{\gamma}{2-\gamma}} < c^{-1} = \left(\frac{f_H}{f_L}\right)_N,$$

In the AI-substitute case, for any $\beta > 0$,

$$\left(\frac{f_H}{f_L}\right)_S = \frac{\gamma^2\beta/c + (1-\gamma)\delta}{\gamma^2\beta + (1-\gamma)\delta} < \frac{\gamma^2\beta/c}{\gamma^2\beta} = \frac{1}{c} = \left(\frac{f_H}{f_L}\right)_N,$$

Evaluating at $\beta_S^* > 0$ preserves the inequality.