



Marketing Science Institute Working Paper Series 2026

Report No. 26-133

Marketing AI Exposure: Applicability Breadth, Substitution Depth, and Job Outcomes

Jia (Jocelyn) Cao, Sunda Bharadwaj, Pengyuan Wang and Shanjukta Nath

“Working Paper Title” © 2026

Authors

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.

Marketing AI Exposure: Applicability Breadth, Substitution Depth, and Job Outcomes

Jia (Jocelyn) Cao, Doctoral Student,
Terry School of Business, University of Georgia, USA,
jocelyn.cao@uga.edu

Sundar Bharadwaj*,
The Coca-Cola Company Chair of Marketing and Professor,
Terry School of Business, University of Georgia, USA,
sundar@uga.edu

Pengyuan Wang, Associate Professor,
Terry School of Business, University of Georgia, USA,
pengyuan@uga.edu

Shanjukta Nath, Assistant Professor
Department of Agricultural and Applied Economics, University of Georgia
shanjukta.nath@uga.edu

* Corresponding author

Conflict of Interest:

The author(s) declare no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding Statement:

This work is supported by the Marketing Science Institute (research agreement SB-4705).

Ethical Approval and Informed Consent Statements:

All validation studies procedures were approved by the Institutional Review Board of the first author's institution at the time of data collection. All participants provided informed consent before participating in the lab or online studies.

Marketing AI Exposure: Applicability Breadth, Substitution Depth, and Job Outcomes

As AI reshapes organizational work, its overlap with marketing responsibilities raises questions about its impact on marketing jobs. This research introduces Marketing AI Exposure as a two-dimensional construct comprising *AI applicability breadth* and *AI substitution depth*.

Applicability breadth captures the range of AI technologies relevant to a job's task portfolio, whereas substitution depth captures the extent to which AI can perform those tasks. Using dependency parsing to extract verb-noun pairs from approximately 4.86 million marketing job postings and 1.3 million AI patents, we construct job posting-level measures of both dimensions. Regression analyses show that breadth is associated with higher hiring but lower wages, whereas depth is associated with lower hiring and lower wages. Jobs high on both dimensions show the strongest wage compression, but not uniform hiring contraction, underscoring the value of distinguishing AI availability from task coverage. The research also shows that marketing skills moderate these relationships heterogeneously.

Keywords: AI Exposure, Automation, Augmentation, Marketing Employees, Marketing Skills, Job Posting Data.

AI is increasingly becoming a critical operational infrastructure for firms, determining how work is organized, allocated, and valued. Yet AI's impact on the labor-market remains unsettled. Early discussions surrounding generative AI emphasized widespread automation and substantial displacement of white-collar workers (e.g., Acemoglu and Restrepo 2019; Cutter and Zimmerman 2025; Seetharaman and Wells 2023). More recent evidence, however, paints a considerably more nuanced picture. Firms increasingly report productivity gains from AI adoption while simultaneously redesigning workflows, expanding hiring in some functions, slowing hiring in others, and reallocating work between humans and AI rather than uniformly substituting workers (Bindley 2026)). Labor economists likewise disagree about AI's employment effects (Anders 2026). While some argue that AI will eventually automate a substantial share of cognitive work, others emphasize productivity gains, new complementary tasks, and organizational adaptation (Nath et al. 2026). Increasingly, the debate has shifted from whether AI will transform work to identifying which tasks AI can perform, where human capabilities remain complementary, and how firms reorganize production around new technologies (Bindley 2026). Together, these developments suggest that AI's labor-market effects are unlikely to resolve into either uniform displacement or universal augmentation. Instead, understanding AI's impact requires identifying the mechanisms through which AI transforms multiple tasks.

This ambiguity extends to the task-based theoretical framework, which suggests that AI's impact shows that it substantially improves performance on some tasks while remaining poorly suited to others, even within the same job (Dell'Acqua et al. 2026). Most jobs comprise a bundle of heterogeneous tasks, such task-level performance differences do not imply uniform employment outcomes for all workers. Firms may therefore respond to AI unevenly across jobs

by expanding output, redesigning workflows, changing the composition of work, or substituting away from selected tasks. Consequently, understanding AI's labor-market effects requires moving beyond occupation-level measures to examine how AI transforms the task composition of jobs.

AI is unlikely to eliminate a job as a whole; instead, it is more likely to automate some tasks within a job while augmenting others. Marketing provides a particularly useful setting for examining this task-level heterogeneity because marketing work is data-intensive, modular, and feedback-rich. Further, it entails combining structured and unstructured data to infer heterogeneous preferences to personalize communication, organizing around separable but interdependent activities, and generating feedback outcomes quickly (Agrawal, Gans, and Goldfarb 2018; Berger et al. 2020; Humphreys and Wang 2018; Netzer et al. 2012; Wedel and Kannan 2016). Recent evidence further indicates that AI can provide accurate, on-demand access to foundational marketing knowledge across a broad range of topics and can increasingly apply that knowledge to marketing problems. Its performance remains less consistent on tasks requiring higher-order reasoning, however, reinforcing the view that AI capabilities vary across the cognitive demands embedded in marketing work (Burke, Matthe, and Leykin 2025). Together, these features and differences in AI capability mean that marketing tasks are unevenly amenable to AI, with some readily automated and others continuing to require human judgment that AI primarily supports.

This task-level perspective raises a broader empirical question. If AI affects different tasks within jobs to different degrees, then why do some jobs experience hiring expansion while others experience hiring contraction? Why do some jobs experience wage compression while others appear relatively protected? Existing measures of AI exposure often summarize these

differences into a single index, making it difficult to distinguish between the range of AI technologies each job can be affected by with the degree to which AI can substitute for a job's tasks. Yet these represent distinct mechanisms with potentially different implications for labor-market outcomes.

We develop *Marketing AI Exposure* as a two-dimensional construct that captures the extent to which marketing jobs are amenable to AI-driven task transformation. The first dimension, AI applicability breadth, refers to the range of AI technologies applicable to a job and reflects the maturity and availability of AI capabilities relevant to that job. The second dimension, AI substitution depth, refers to the extent to which AI can perform the job's tasks. Together, breadth and depth capture the conditions under which AI adoption becomes more likely. Breadth increases the feasibility of adoption by expanding the set of applicable technologies, whereas depth increases the value of adoption by expanding the share of job tasks AI can perform. These dimensions need not move together. A job may score high on breadth because many AI technologies are relevant, yet low on depth if those technologies affect only a narrow portion of its task bundle. Conversely, a job may score low on breadth but high on depth if a smaller set of AI tools can perform many of its core tasks. A unidimensional measure would conflate these cases and obscure whether marketing jobs face broad technological applicability, deep task substitution, or both. We therefore examine two research questions: (1) how AI applicability breadth and AI substitution depth affect wages and hiring in marketing jobs, and (2) how demand for different marketing-specialized skills moderates these relationships.

We measure the AI marketing exposure using a verb-noun pair extraction methodology applied to more than 4.86 million marketing job postings from Lightcast from 2020 to 2024 and match these postings to AI patent records from the USPTO. The regression results show that

applicability breadth is associated with higher hiring but lower posted wages, while substitution depth is associated with both lower posted wages and reduced hiring demand. Jobs high on both breadth and depth show the strongest wage compression, but not uniform hiring contraction relative to low-breadth, low-depth jobs. This pattern underscores the value of a two-dimensional measure: breadth and depth identify distinct mechanisms that a single AI-exposure index would not be able to disentangle.

The post hoc moderator analysis further shows that AI breadth and depth reprices marketing skills¹ rather than uniformly raising or lowering their value. AI skills, brand management, and marketing operations skills attenuate wage penalties associated with applicability breadth but amplify wage penalties associated with substitution depth, suggesting that these skills retain value when AI broadens what workers can do, but become less protective when AI can perform the underlying work. These wage effects, moreover, do not always move in step with hiring effects. Customer management, sales, and integrated marketing communication skills more consistently offset wage penalties than hiring contraction, while marketing operations is the most consistent hiring-protective category because firms still need workers to integrate AI into campaign systems, dashboards, governance routines, and cross-functional workflows. Overall, the moderator results show that AI reallocates demand across marketing task bundles and reprices the skills attached to those bundles.

The paper makes three contributions. First, it brings marketing human capital into the AI labor-market debate by examining the occupations responsible for performing marketing work, an area that has received limited attention in marketing research. Second, it distinguishes two dimensions often conflated in AI-exposure measures: the availability of AI technologies

¹ Please see web appendix G for skill classification

applicable to a job and the extent to which AI can perform its core task bundle. Their different associations with posted wages and hiring demonstrate the empirical importance of this distinction. Broad but shallow applicability is associated with higher hiring and lower wages, whereas greater substitution depth is associated with lower wages and hiring. When both dimensions are high, wage compression is especially pronounced, but hiring is not uniformly lower. Third, the post-hoc moderator analysis shows that marketing-specialized skills condition these relationships differently across wages and hiring. Skills centered on codifiable execution tend to be associated with greater labor-market pressure, whereas skills involving customer interaction, coordination, and judgment more often attenuate negative associations. These findings provide a more differentiated account of how marketing skills relate to labor-market outcomes as AI becomes applicable to a wider range of tasks and capable of performing more core work.

Relevant Literature on AI and Human Capital

Three streams of literature are relevant to this study and Table 1 provides a summary of the literature. First, research on marketing human capital shows that marketing talent matters for firm performance, centered largely on CMOs, salespeople, and salesforce organizations (e.g., Ahearne et al. 2008; Bansal et al. 2017; Germann, Ebbes, and Grewal 2015; Lawrence et al. 2021; Srivastava, Kashmiri, and Mahajan 2023; Whitler, Krause, and Lehmann 2018). This literature establishes the importance of marketing human capital, but it devotes little attention to the broader set of marketing roles and how AI affects marketing work.

In parallel, marketing research on AI has emphasized algorithmic applications to marketing problems, consumer responses to AI, and firm-level AI strategy and implementation. The first applies AI methods to core marketing problems such as targeting and sponsored search

(e.g., Kaiser and Schulze 2026; Lu and Kannan 2025; Wang, Zhang, and Zhang 2026). A second examines customer responses to AI agents, AI advice, and AI-mediated service, highlighting the roles of competence, warmth, authenticity, and task type (e.g., Garvey, Kim, and Duhachek 2023; Li, Lai, and Wang 2025; Longoni and Cian 2022; Tully, Longoni, and Appel 2025). A third focuses on firm-level AI strategy and implementation (e.g., Hartmann, Exner, and Domdey 2025; Kapoor and Kumar 2025; Mu and Zhang 2025). Collectively, these studies say relatively little about how AI reshapes the tasks that constitute a marketing job itself, and thus the marketing employees that perform it.

Another stream links AI to various labor-market outcomes. At the firm level, AI investment has been associated with productivity, employment growth, sales, and market valuation (Babina et al. 2024; Damioli et al. 2023; Eisfeldt et al. 2023). At the occupational level, some studies find that greater AI exposure reduces labor demand, especially when AI substitutes for human work (Acemoglu et al. 2022; Hampole et al. 2025; Hui, Reshef, and Zhou 2024), whereas others emphasize complementarity, productivity gains, and changes in skill composition (Brynjolfsson, Chandar, and Chen 2025; Nath et al. 2026). Concurrent marketing studies examine ChatGPT's release as an exogenous shock and find its effect on marketing positions and marketing job postings uneven across experience and task levels (Luo, Miao, and Sudhir 2026; Zhu et al. 2026). However, they do not directly measure how a job's task portfolio shapes its exposure to AI.

The third stream focuses on AI occupation-level exposure measurement. Labor-economics research has developed measures of occupational AI exposure, typically by mapping AI capabilities to task descriptions in O*NET and related sources (Brynjolfsson, Mitchell, and Rock 2018; Eloundou et al. 2024; Felten, Raj, and Seamans 2021; Webb 2019). However, most

Table 1. Empirical Literature on AI and Job Outcomes

Study	AI Measure	Focal Dependent Variable(s)	Moderators	Marketing Focused	Level of Analysis	Main Findings
Felten, Raj, and Seamans (2021)	AI Occupational Exposure (AIOE)	-	-	No	Occupation	Create a new measure of AI occupational exposure and assess AI exposure measure at industry and geographic levels.
Acemoglu et al. (2022)	AI exposure	Job vacancies; skill requirements	-	No	Occupation	Rapid growth in AI related vacancies; AI-exposed establishments reduced non-AI hiring and induced significant negative and positive skill changes
Babina et al. (2024)	AI investment	Productivity; employment growth	-	No	Firm	Firms that invest more in AI experience higher productivity and employment growth
Eloundou et al. (2024)	LLM exposure	Wage; productivity ²	-	No	Occupation	Higher-wage occupations are facing greater exposure to LLMs; productivity growth reveals little relationship with exposure to LLMs.
Hui, Reshef, and Zhou (2024)	Gen AI models introduction	Freelancer employment; earnings	Freelancer's past performance	No	Occupation	Freelancers who offer service in high AI-exposed occupations experienced reduction in both employment and earnings
Luo, Miao, and Sudhir (2026)	ChatGPT emergence	Marketing employment; salary	Seniority; Resolvability (production vs. orchestration)	Yes	Firm	Generative AI's emergence unevenly affects marketing employment, with effects varying by worker experience and task types; Salary rise across all job functions
Zhu et al. (2026)	ChatGPT emergence	Marketing employment	Seniority; job types	Yes	Firm	Marketing headcount declined by 0.92% relative to non-marketing peers; the relative decline varied by seniority levels and job types
This study	Two-dimensional measure of AI Application Breadth and Substitution Depth	Job Salary and hiring	Marketing Skills	Yes	Job	Breadth increases hiring but lowers wages; depth lowers both wages and hiring; joint presence is most detrimental; marketing skills heterogeneously moderate these effects

² In Eloundou et al. (2024), wages and productivity are predictors rather than dependent variables.

measures are effectively unidimensional, collapsing all the variation into one single index: how many AI technologies are relevant to a job and what share of its tasks those technologies can actually perform. For example, a job can score high on either dimension alone, or existing measures cannot tell these cases apart. This distinction is especially important in marketing, where many jobs may be affected by AI but not equally replaceable by it.

Taken together, these streams reveal a clear opportunity. Research has not developed a multidimensional measure of AI exposure that distinguishes the breadth of AI applicability from the depth of AI substitution, nor has it linked those distinct dimensions to marketing job wages, and postings, nor examined the moderating role of marketing skills.

Conceptual Background

Marketing Work and AI Capability Alignment

A task-based perspective explains why AI may complement some elements of marketing work while substituting for others. Technology affects the tasks within occupations rather than necessarily replacing occupations as a whole (Autor, Levy, and Murnane 2003). Marketing is particularly relevant because many of its tasks align with AI capabilities in prediction, generation, and optimization (Agrawal, Gans, and Goldfarb 2018; Davenport et al. 2020; Sutton and Barto 2018). This alignment, however, varies with the data, workflow structure, and feedback associated with each task.

First, marketing relies extensively on structured and unstructured data for forecasting, customer analysis, personalization, media allocation, and communication (Lemmens et al. 2025; Singh, Sen, and Borle 2022; Wedel and Kannan 2016). Tasks based on codifiable, high-volume data are more amenable to predictive and generative AI than tasks requiring tacit knowledge, contextual interpretation, relational judgment, or accountability for market outcomes

(Brynjolfsson, Mitchell, and Rock 2018; Eloundou et al. 2024).

Second, marketing activities can often be separated into discrete but interdependent workflow components. This structure allows AI to perform specific activities such as lead scoring, reporting, campaign execution, or media allocation. Their interdependence nevertheless preserves a role for human coordination and strategic choice when AI-performed activities must be integrated with other elements of marketing work (Davenport et al. 2020).

Third, many marketing activities generate observable response measures, including engagement, conversion, retention, and sales. Such signals support AI-based optimization when they are timely and attributable to particular actions (Agrawal, Gans, and Goldfarb 2018; Kumar, Ashraf, and Nadeem 2024; Sutton and Barto 2018). AI is less reliable when outcomes are delayed, noisy, or difficult to attribute, leaving greater scope for human interpretation.

Marketing jobs therefore combine tasks with different levels of amenability to AI. Codifiable and repeatable cognitive tasks are generally more susceptible to automation, whereas creative reasoning, relational judgment, and strategic integration remain more reliant on human capabilities (Autor and Dorn 2013; Eloundou et al. 2024; Felten, Raj, and Seamans 2021). Because these tasks coexist within jobs, AI may substitute for some activities while complementing others. Measuring AI exposure at the job level consequently requires accounting for its alignment with the underlying task bundle's composition.

A Two-Dimensional View of Marketing AI Exposure

Task-technology fit theory provides the conceptual foundation for understanding how AI capabilities align with tasks that constitute marketing work. Technology use and performance depend on the fit between task requirements and technology capabilities (Furneaux 2012; Goodhue and Thompson 1995). Applied to AI, technology adoption potential depends on

whether AI capabilities fit the tasks that define a marketing job. High fit increases the feasibility of AI execution, whereas low fit preserves human involvement. Fit, however, is not a single dimension quantity. It depends jointly on the maturity of the AI technology ecosystem and the extent to which AI capabilities cover that job's core tasks, and these two quantities need not move together (Howard and Hair 2023; Mahmood, Hall, and Swanberg 2001). A job's technology ecosystem can expand without that expansion reaching a large share of the job's tasks, just as a small number of tools can reach deep into a job's important tasks. A single dimension fit measure would average over this difference; therefore, we decompose fit into two determinants: how many AI technologies could potentially apply to a job and how much of the job's actual task content those AI technologies execute.

AI applicability breadth is the extent to which a job is associated with a wide range of AI technologies that can support or perform its task. High breadth indicates a more mature technology ecosystem, lower adoption uncertainty, and greater availability of deployable tools. AI substitution depth is the extent to which AI can execute a large, importance-weighted share of a job's task bundle (Arntz, Gregory, and Zierahn 2017; Eloundou et al. 2024). High depth indicates that AI reaches a larger portion of the job's important tasks. This distinction is conceptually important. Jobs with many relevant AI tools may appear highly exposed even when those tools mainly assist human workers. Conversely, jobs with fewer AI systems may face substantial substitution risks if those systems impact important tasks. A unidimensional exposure measure would conflate AI availability with AI replaceability.

Labor-Market Consequences for Marketing Occupations

AI exposure can affect wages and hiring through substitution or complementarity. When AI performs tasks previously undertaken by workers, firms may reduce labor demand and wages.

When it complements workers, productivity gains may increase marketing activity and hiring (Acemoglu et al. 2022; Hampole et al. 2025; Luo, Miao, and Sudhir 2026). Increased hiring, however, need not protect wages. AI may enable firms to undertake more work while standardizing execution and reducing the scarcity value of some human inputs. We therefore distinguish labor-expanding augmentation, which increases hiring, from wage-enhancing augmentation, which raises returns to worker expertise, judgment, or coordination capabilities (Brynjolfsson, Li, and Raymond 2025; Nath et al. 2026).

Greater applicability breadth makes adoption more feasible by giving firms a wider range of technologies to support marketing work. This availability may increase hiring as firms undertake more campaigns, analyses, experiments, and customer interactions. It may simultaneously lower wages if widely available tools standardize execution, reduce dependence on specialized skills, or expand the pool of workers capable of performing the job. Breadth is therefore more likely to produce labor-expanding augmentation, which increases hiring while reducing the value of some human inputs.

Substitution depth indicates whether AI reaches peripheral activities or important job responsibilities. Because deep AI capability displaces rather than merely assists on those tasks, higher depth should reduce both hiring and wages. Firms need fewer workers to complete the same tasks and the employees lose bargaining power over tasks AI can perform directly. When breadth is low and depth is high, these negative effects on hiring and wages may therefore remain limited since a narrower technology ecosystem limit how readily firms will adopt and deploy AI even where its capability is deep.

When both dimensions are high, AI is widely available and capable of performing a substantial share of core responsibilities. Depth should consequently weaken the positive

relationship between breadth and hiring and intensify wage pressure. Broad applicability may still increase the amount of marketing work undertaken, even as depth reduces the human labor needed for parts of the task bundle. The two dimensions jointly determine whether AI increases hiring, places downward pressure on wages, or reduces both.

The distinction is consequential in marketing because jobs vary sharply in task codifiability, required judgment, and dependence on market context. Reporting, ad trafficking, campaign execution, and some marketing analytics involve structured and repeatable tasks that AI can perform or accelerate. Account management, business development, product marketing, and brand strategy rely more heavily on persuasion, organizational judgment, creative direction, and responsibility for market outcomes. AI may improve information processing and execution in these roles, but their value remains tied to human judgment and domain-specific expertise.

The effects of AI exposure should therefore depend partly on the marketing-specialized skills embedded in a job. These skills are disproportionately concentrated in marketing roles and distinguish marketing professionals from adjacent labor pools. Skills centered on codifiable execution may amplify wage or hiring pressure because AI can standardize, automate, or unbundle the associated tasks. Skills involving integration, coordination, strategy, and contextual judgment may mitigate these effects by complementing AI-enabled work. Marketing skills thus moderate the relationships among applicability breadth, substitution depth, wages, and hiring.

Empirical Setting

Data and Sample

We combine online job-posting data with AI patents to measure two dimensions. Job postings capture downstream task content and labor-market demand for marketing jobs, while AI patents serve as a proxy for the upstream technological supply underlying AI capabilities, reflecting technological innovation and maturity (Fleming and Sorenson 2004). The job-posting

data is sourced from Lightcast (formerly Burning Glass), which aggregates online job vacancies across different sources (e.g., job boards and employer websites), and standardizes them after removing duplicates. Each posting includes information on the employer, job title, posting date, location, salary, education and experience requirements, skills, occupation codes, and full job text. The granted patent data is from the Artificial Intelligence Patent Dataset, which classifies USPTO patents as AI-related or non-AI-related and provides the patent title, publication date, and abstract. We use patent abstracts to characterize AI capabilities, because abstracts summarize a technology's functions without the legal and procedural language, and use publication dates to ensure that, for each posting, exposure is based only on AI technologies available before the posting date.

The Lightcast data includes more than 200 million U.S. job postings from January 2020 through October 2024 and each posting is classified into one occupation group. Because the empirical object is marketing work rather than all labor-market activity, we identify marketing-relevant occupations through a structured expert procedure. Six marketing professionals evaluate candidate occupation groups, and we retain groups classified as marketing-related by the majority of respondents (see web appendix A for details). Tied cases are resolved through author review.

We begin with 23 million marketing job postings. We then remove postings from staffing and temporary employment agencies because these postings often obscure the actual hiring firm, and we exclude internship and part-time postings because their task structures differ from full-time roles. This procedure yields 24 broad marketing-related occupation groups, 353 sub-occupation groups, and more than 730,000 job titles, resulting in a sample of more than 16 million full-time marketing job postings. Our salary analyses use the subset of postings that

report salary information, totaling 4,859,324.

Measuring AI Applicability Breadth and AI Substitution Depth


Our measurement task is to capture the overlap between the task requirements of marketing employees and the potential capabilities of AI technologies. Both job postings and patent abstracts are unstructured texts, so direct comparison requires a common representation of tasks and technological capabilities. We use dependency parsing to extract verb-noun pairs from both sources. In job postings, these pairs represent the actions and objects that define marketing tasks. In patent abstracts, they represent the functions that AI technologies can perform. For example, the job task “create comprehensive multimedia communication plans” yields the pair “create + plan.” A patent abstract that describes a system for coordinating customer activities yields “coordinate + activity.” These pairs provide the common action-object unit for matching marketing tasks to AI capabilities.

The measurement procedure has three steps, as shown in Figure 1. First, we parse job responsibilities and AI patent abstracts, lemmatize terms, extract verb-direct object pairs, and recover prepositional objects when the relevant object occurs within a prepositional phrase. We then map nouns to WordNet hypernyms to reduce dependence on surface vocabulary. For instance, “clients” and “users” can both be represented as “customer,” allowing semantically equivalent tasks to match even when different words are used. Based on a human validation study comparing WordNet hierarchy levels, we use level 6 in the main analysis. Details of parsing, WordNet mapping, and hierarchy selection are reported in the web appendix B.

Next, we match verb-noun pairs across job tasks and AI patents. A job-task pair is treated as AI-executable if the same pair appears in at least one AI patent published before the job posting date. This is a conservative matching procedure because it requires equivalence at the

action-object level rather than at a broad semantic level. The WordNet normalization applied in Step 1 mitigates this concern by mapping synonymous terms to shared abstract categories before matching. To illustrate, consider the following job task and patent abstract:

Figure 1. Empirical Approach of Artificial Intelligence Marketing Replacement



Task A: “*Coordinate Joint Go-to-Market activities.*”

Patent A: “A system, method and computer program product for integrating, managing and *coordinating activities* of customers in an enterprise, across a plurality of applications, resources and touchpoints.”

Because dependency parsing extracts “coordinate + activity” from both sources and “activities” are normalized to their lemma “activity” in Step 1, patent A is matched to task A. This match indicates that the function encoded in Patent A directly corresponds to an action that marketing employees must perform. Notably, this matching is many-to-many, meaning a single task may be matched to multiple patents and a single patent may be matched to multiple tasks (which forms the basis of our exposure measure).

Then, we construct two measures. AI substitution depth captures the extent to which the task content of a marketing job can be performed by AI technologies, measured as the proportion of verb-noun pairs within a job that are matched to at least one AI patent. At the task-level, each task is presented by a set of verb-noun pairs extracted from its description. At the pair level, we

define a binary substitution depth indicator, a_{ij} , coded as 1 if the verb-noun pair exactly matches at least one verb-noun pair in the AI patent database, and 0 otherwise. At the job level, AI substitution depth is aggregated across all verb-noun pairs within a job, with each task weighted by its O*NET importance score w_i to reflect the centrality of that task to the job. Formally:

$$Substitution\ Depth_o = \frac{\sum_{i=1}^I \sum_{j=1}^{J_i} a_{ij} w_i}{\sum_{i=1}^I \sum_{j=1}^{J_i} w_i},$$

where o denotes the job, i indexes tasks 1 through I within the job o , j indexes verb-noun pairs 1 through J_i within task i , a_{ij} is the binary substitution depth indicator for pair j of task i , and w_i is the importance score of task i from O*NET importance score of task i (see below for details on how w_i is constructed).

AI applicability breadth captures how many distinct AI technologies are potentially applicable to a job. At the pair level, breadth is defined as the number of patents whose verb-noun pairs match a given verb-noun pair extracted from each task description of the job. At the job level, breadth is aggregated across all verb-noun pairs within a job and standardized by the total number of AI patents published prior to the job posting date. Dividing by P standardizes exposure across job postings from different time periods, ensuring that the measure reflects the relative prevalence of AI technologies applicable to a task rather than the absolute growth of the patent stock over time.

$$Applicability\ Breadth_o = \frac{1}{P} \frac{\sum_{i=1}^I \sum_{j=1}^{J_i} e_{ij} w_i}{\sum_{i=1}^I \sum_{j=1}^{J_i} w_i},$$

where o denotes the job, i indexes tasks 1 through I within the job o , j indexes verb-noun pairs 1 through J_i within task i , e_{ij} is the number of AI patents whose verb-noun pairs match pair j of task i , and P is the total number of AI patents published prior to the date of job posting, w_i is the importance score of task i from O*NET.

Higher applicability breadth values indicate that a greater variety of AI technologies are potentially relevant to the job at the time of posting. Together, depth captures task coverage, while breadth captures the available options of the relevant AI technology ecosystem.

Both measures are computed as an importance-weighted proportion of verb-noun pairs matched at the pair level, rather than an average of task-level match rates. That is, tasks are not only weighted by their O*NET importance score (w_i) but also contribute proportionally to their number of decomposed verb-noun pairs (J_i), such that tasks described with greater granularity in a posting (i.e., yielding more identifiable pairs) receive correspondingly greater weight in the aggregate measure.

Task importance weights come from *ONET* and details are provided in web appendix C (Eloundou et al. 2024; Felten, Raj, and Seamans 2021). Because job-posting task descriptions do not always exactly match ONET task descriptions, we link each posting task to the closest ONET task within the same occupation using cosine similarity from a sentence-embedding model and assign the corresponding ONET importance score (Hampole et al. 2025). This procedure gives greater weight to tasks that are central to the occupation rather than treating all responsibilities as equally informative. Please see web appendix D for the real-case example pipeline.

Measure Validation

We conducted five validation exercises and provide full validation procedures, stimuli, participant exclusions, and statistical tests in the web appendix E. First, we validated the verb-noun extraction procedure with a Prolific study. Participants evaluated whether extracted verbs, nouns, and verb-noun pairs accurately represented the original task statements. The results support the use of verb-noun pairs as a compact representation of task content and support WordNet level 6 (see web appendix B for details on abstraction level) as the preferred level of

abstraction. Second, we recruited participants from Prolific to validate whether the verb-noun pair matching method captures the potential applicability of AI technology to a given task. We found the correct classification rate is 72.59%, suggesting that the verb-noun pair approach reasonably captures human perception of AI's capability in marketing tasks. Third, we conducted a 2×2 experiment that manipulated breadth and depth in a marketing job description. The results indicate that the participants distinguished between the two dimensions in manipulation checks. Fourth, we asked experienced marketing practitioners whether the conceptual distinction between AI applicability breadth and AI substitution depth aligns with their understanding of AI's effects on marketing work. Practitioners recognized that a job can involve many relevant AI technologies without those technologies covering the important task bundle, and that a smaller set of AI technologies can create greater replacement pressure when they perform a large share of the work. These results support the conceptual claim that breadth and depth are related yet distinct aspects of marketing AI exposure and demonstrate the case for a two-dimensional exposure measurement. Finally, to distinguish our measure from the established Eloundou et al. (2024) measure, we examined the correlations between our measures and found them to be low: breadth correlated at -0.055, and depth correlated at -0.215.

Descriptive Patterns

As reported below in Table 2³, the salary subsample is broadly similar to the full marketing-posting sample on the two exposure measures ruling out the likelihood that the two data sets vary in AI exposure. Mean applicability breadth is nearly identical across the full sample and salary subsample; mean substitution depth is also similar; the effect size for both is negligible (Cohen's d smaller than 0.03).

³ Breadth is the count of distinct matched AI patent verb-noun pairs divided by the total number of AI patents published before the posting date, yielding super small values because AI patents exceed one million.

Table 2. Summary of Key Measure (Full Sample vs. Focused Sample)


	Raw Lightcast Dataset		Salary Sample	
	Mean	SD	Mean	SD
Applicability Breadth	.0003	.0004	.0003	.0004
Substitution Depth	.740	.242	.744	.246
Observations	16,775,310		4,859,324	

Table 3 presents the distribution of marketing/sales-related job postings across occupation groups from 2020 to 2024. As shown in Figure 2, the top 5 marketing job vacancies increased substantially over the sample period. Across the 24 occupation groups, job postings rose from roughly 400,000 in 2020 to more than 1.2 million in 2024. This increase is distributed across multiple occupation groups rather than concentrated in a single role, indicating broad growth in marketing-related online vacancies during the period.

Table 3. Occupation Group Distribution Over Time


	2020	2021	2022	2023	2024
Non-Technical Sales	121,445	171,711	191,118	210,683	213,950
Customer Service Representatives	82,463	138,519	168,371	168,096	165,928
Retail Sales and Service	47,967	93,308	124,852	155,420	171,307
Retail Managers	30,850	60,067	89,533	131,090	146,125
Insurance Claims and Sales	27,217	46,263	50,810	48,721	45,482
Account Managers and Executives	26,323	42,716	59,136	77,706	79,926
Sales Managers	17,275	29,287	43,115	66,205	72,388
Marketing Specialists	16,302	26,728	33,265	37,439	34,636
Merchandising	15,203	26,854	38,668	46,588	46,090
Customer Service Managers	10,215	19,538	30,966	33,962	35,775
Business Analysis	9,265	23,493	43,154	48,506	48,859
Communications and Public Relations	9,202	19,075	29,711	35,791	34,787
Visual Design / Arts	7,309	10,984	15,810	18,176	17,207
Marketing Managers	6,818	21,077	35,764	50,640	55,980
Hospitality and Travel	5,831	10,566	18,027	19,912	18,833
Data Analysis and Mathematics	5,236	13,336	29,539	37,862	38,484
Energy Sales	3,515	4,590	5,162	6,001	5,184
Technical Sales	3,364	4,479	6,522	9,749	9,644
Fundraising and Grants	3,210	6,799	11,684	13,841	18,819
Business Intelligence	3,140	7,088	14,911	18,405	19,720
Financial Sales	3,032	6,806	10,187	15,662	17,776
Hospitality Managers	1,266	3,505	5,929	6,460	6,205
Writing and Language	401	832	1,381	1,190	1,163
Promotional / Commercial Modeling	36	87	128	165	150

Figure 2. Occupation Group Distribution Over Time (Top 5)



The two exposure measures are positively related but far from redundant. The correlation between AI applicability breadth and AI substitution depth is .256, indicating that the two measures share only 6.55% of their variance and thus capture largely independent aspects of AI's applicability to marketing roles. The descriptive analysis supports the value of decomposing AI exposure into breadth and depth. Figure 3, Panel A, shows that marketing jobs are distributed across all four breadth–depth profiles, including substantial off-diagonal cases. These off-diagonal cases are theoretically important because they distinguish jobs broadly affected by AI from jobs that are deeply substitutable by a narrower set of technologies. Some roles, such as claims specialist and retail sales, have high depth but low breadth, meaning that a limited set of AI technologies covers a large share of the task bundle. Others, such as consulting, firm sales, and merchandising managers, have high breadth but low depth, meaning that many AI technologies are relevant but do not penetrate deeply into the core task structure. These off-diagonal jobs are precisely the cases that a unidimensional exposure measure would obscure. See web appendix F for AI breadth and depth variation across occupation groups and industries.

Figure 3. Descriptive Statistics



Econometric Analysis

Outcomes and Moderators

We examine two labor-market outcomes: salary and hiring. Salary is the advertised salary in the job posting. When a posting reports a range, Lightcast provides the midpoint, which we use as the salary measure. Hiring is measured by the number of job postings for the same sub-occupation title, firm, and MSA location in the same quarter. We use job postings as a measure of labor demand, consistent with prior research using online vacancies. The salary analysis is conducted at the posting level. The hiring analysis aggregates and averages outcomes, exposure and controls at the firm-sub-occupation-location-quarter level.

We also examine whether skill requirements moderate the effects of AI exposure. Lightcast classifies skills using machine learning and expert-validated taxonomies. We construct skill-share measures as the ratio of each skill category to the total number of skills listed in the

posting. The focal categories are marketing analytics, AI skills, customer management skills, sales skills, brand management skills, marketing operations skills, and integrated marketing communication (IMC) skills (see web appendix G for details). These categories allow us to distinguish skills that are more technical, customer-facing, analytical, requiring coordination, or communication-oriented.

To account for alternative factors driving salary and job hiring outcomes, we include a set of job-posting level control variables (see Table 4). Specifically, we account for the number of responsibilities listed, minimum year of experiences required, and education required that are included in each posting (Nath et al. 2026), and these control variables are averaged in job-hiring analysis. In addition to job-posting level control variables, we also include the presence or absence of laws requiring salary disclosure, which has been shown to independently affect wage levels and the availability of salary information in job postings (Arnold, Quach, and Taska 2025). Laws requiring salary disclosure is a binary indicator equal to one if the job posting location has an effective law requiring employers to disclose salary ranges for open positions. Although recent research finds no significant relationship between salary disclosure laws and the number of job postings (Arnold, Quach, and Taska 2025), we include this control variable in the job hiring analysis.

Table 4. Key Variable Operationalizations and Data Sources

Variables	Data Source	Operationalization
Dependent Variables		
Salary	Lightcast	The advertised salary reported in the job posting
Job Hiring	Lightcast	The total number of job postings with the same sub-occupation title published by the same company in one quarter for the same location
Skill Requirements		
Skills	Lightcast	The ratio of AI/Sales/Customer-management/Brand Management/Marketing Operations/Marketing Analytics/IMC skills and the total number of skills
Independent Variable		

Depth	Lightcast, AI patents	The depth score of each job posting
Breadth	Lightcast, AI patents	The breadth score of each job posting
Control Variables		
Number of Responsibilities	Lightcast	The total number of responsibilities listed in the job posting
Minimum years of experience	Lightcast	Number of years of experience required in the job posting
Education required	Lightcast	Binary indicator equal to one if the job posting requires any level of education
State law for salary disclosure	GovDoc	Binary indicator equal to one if the job posting location has an effective law of salary disclosure

Baseline Specification

The salary specification estimates the association between AI exposure and posted wages while conditioning on job, firm, occupation, geography, and time characteristics:

$$Salary_{ifomt} = \beta_1 Depth_{ifomt} + \beta_2 Breadth_{ifomt} + \beta_3 (Depth_{ifomt} \times Breadth_{ifomt}) + X_{ifomt}\theta + \alpha_f + \mu_o + \gamma_m + \lambda_t + \epsilon_{ifomt}, \quad (1)$$

The dependent variable is the standardized salary for job posting (i), firm (f), occupation (o), market (m), and time (t). The focal regressors are standardized AI substitution depth, standardized AI applicability breadth, and their interaction. The interaction tests the labor-market effects when AI technologies are both widely available and capable of performing a greater share of the task bundle. Controls include education requirements, experience requirements, number of responsibilities, skill shares, salary-disclosure-law indicator, and two control function residuals. Fixed effects include firm, occupation, MSA, and year-month.

The hiring specification estimates hiring at the firm-sub-occupation-location-quarter level:

$$Job\ Hiring_{sfomt+1} = \beta_1 AvgDepth_{sfomt} + \beta_2 AvgBreadth_{sfomt} + \beta_3 (AvgDepth_{sfomt} \times AvgBreadth_{sfomt}) + X_{sfomt}\theta + \alpha_f + \mu_o + \gamma_m + \lambda_t + \epsilon_{sfomt}, \quad (2)$$

where $Job\ Hiring_{sfomt+1}$ is the total number of job postings for sub-occupation groups at the

firm j at time $t + 1$. $Depth_{sfomt}$ and $Breadth_{sfomt}$ are average AI depth and breadth scores for the title i at firm f at time t . Fixed effects include α_f (firm), μ_o (occupation), γ_m (MSA) and λ_t (year-quarter). The vector X_{sfomt} includes average minimum years of experience requirements, average number of responsibilities, number of job postings that have education requirements, salary-disclosure-law indicator and two control function residuals.

Identification

The empirical design uses within-firm, within-occupation, within-location, and within-time variation in AI exposure. Firm fixed effects absorb time-invariant differences in employer quality, compensation policy, technology orientation, and recruiting strategy. Occupation fixed effects absorb persistent differences across marketing roles. MSA fixed effects account for time invariant local labor-market conditions. Year-time fixed effects absorb aggregate shocks common to all postings, including macroeconomic changes and national shifts in AI salience.

Although firm, occupation, and location fixed-effects absorb time-invariant factors, they do not account for unobservables that are varying with time. We expect AI applicability breadth to be largely exogenous to any single firm, because it reflects the innovation development across the broader AI ecosystem rather than decisions made by individual employers. Even so, firms that are actively investing in AI may simultaneously design the job descriptions and adjust salaries or hiring demand in response to their broader technology strategy. This endogeneity concern motivates our identification analysis. To address the joint endogeneity of AI depth and breadth, and to identify main effects as well as their interaction effect, we employ a two-step control function with two instrumental variables (Wooldridge 2015).

We instrument for applicability breadth and substitution depth using AI patent stock of industry peers and the number of data centers at the MSA-time level. Peer patent stock predicts

the industry-level patent supply of relevant AI technologies but is less likely to be driven by the focal firm's contemporaneous demand for a specific marketing role, making it informative primarily for applicability breadth. The number of data centers in the posting's MSA proxies for local computational infrastructure available to deploy AI into task execution but is similarly unlikely to reflect a single firm's hiring decisions, making it informative primarily for substitution depth. Following the control function, we retrieve the predicted residual $\hat{v}_{1,s/ift}$ and $\hat{v}_{2,s/ift}$, by regressing the endogenous regressors on the two instruments and control variables and include the residuals obtained from this first step in equation (1) and (2). Under the assumption of $E[u|z_1, z_2, breadth, depth] = E[u|v_1, v_2] = p_1 v_1 + p_2 v_2$ (3), where u denotes the structural error term in the original structure equations (1) and (2), the inclusion of two residuals in the second stage should absorb time-varying omitted factors with depth, breadth and their interaction and no separate residual term for interaction term is required. The correction implied by (3) is governed by the assumed joint distribution of the error terms rather than by the functional form in which depth and breadth enter the structural equation, specifically, the interaction term is a function of depth and breadth, its correlation with u is not an independent source of endogeneity. It is therefore already absorbed once two residuals are included. Substituting (3) into the main equation, the remaining error $e = u - p_1 v_1 - p_2 v_2$ satisfies our identification assumption that the remaining error is uncorrelated with our main independent variables when we condition on a set of fixed effects, posting-level covariates, and two extracted first-stage residuals. We assume:

$$E[e_{s/ift}|Depth_{s/ift}, Breadth_{s/ift}, Depth_{s/ift} \times Breadth_{s/ift}, X_{s/ift}, \hat{v}_{1,s/ift}, \hat{v}_{2,s/ift}, \alpha_f, \mu_o, \gamma_m, \lambda_t] = 0$$

Web appendix Tables H1 and H2 report the first-stage control-function estimates.

Because of the weak-instrument concern, we also conduct Anderson–Rubin tests, which yield statistics of 30.07 ($p = .000$) for the salary model and 11.26 ($p = .000$) for the hiring model. Following Keane and Neal (2023), these results provide weak-instrument-robust evidence that the estimated effects are jointly different from zero and are not driven solely by conventional 2SLS inference under potentially weak instruments. Together, the fixed-effects specification accounts for stable heterogeneity at firm, occupation, and location levels as well as common time trends, while the instrument approach addresses time-varying endogenous exposure arising from firm-level decisions.

Results

AI exposure and hiring

Table 5 reports the hiring results (the full table shown in web appendix I1), main variables are standardized and the highest VIF is 3.25 suggesting that multicollinearity is not a concern. As shown in model 1 (main-effect only), AI applicability breadth is positively associated with marketing hiring ($\beta = 1.307$, $p = .000$), whereas AI substitution depth is negatively associated with hiring ($\beta = -2.214$, $p = .000$). When we add the interaction term of breadth and depth, it is negative and significant ($\beta = -0.064$, $p = .000$). These results support the distinction between broad AI applicability and deep AI substitutability. Breadth is associated with more hiring, suggesting that a richer AI technology ecosystem may expand or reorganize marketing work. Depth has the opposite association. When AI covers more of a job's task bundle, firms post fewer jobs. The negative interaction indicates that the positive breadth-hiring association weakens when depth is high. Broad AI availability is therefore not uniformly expansionary. Its labor-demand benefit declines when AI can perform more of the core work. To translate our results into interpretable terms, we estimate the model by using $\log(\text{postings})$ as the

dependent variable and standardized breadth and depth as independent variables (see column 9 of web appendix Table J1). The predicted increase in hiring for high-breadth/low-depth relative to low-breadth/low-depth is approximately 6.5%, approximately 4.1 additional postings per 100 positions (mean = .648)⁴. High breadth and depth show a small net decline (-0.64%) relative to low breadth and depth, indicating the results remain positive throughout the observed range of depth instead of a reversal in the marginal effect.

Table 5. AI Depth, Breadth, and Job Hiring

	Model 1			Model 2		
	Estimate	SE	P-value	Estimate	SE	P-value
Applicability Breadth	1.307	.238	.000	1.448	.240	.000
Substitution Depth	-2.214	.430	.000	-2.329	.431	.000
Breadth × Depth				-.064	.002	.000
Marketing Analytics × Breadth				-.091	.028	.001
Customer-management × Breadth				.007	.007	.285
AI × Breadth				-.134	.061	.027
Sales × Breadth				.012	.005	.011
Brand Management × Breadth				.098	.034	.004
Marketing Operations × Breadth				.100	.013	.000
IMC × Breadth				-.067	.011	.000
Marketing Analytics × Depth				.029	.022	.191
Customer-management × Depth				-.012	.011	.273
AI × Depth				-.051	.032	.114
Sales × Depth				-.056	.004	.000
Brand Management × Depth				-.054	.021	.011
Marketing Operations × Depth				.027	.010	.006
IMC × Depth				-.043	.007	.000
Control Function Residual (Breadth)	-1.296	.238	.000	-1.411	.240	.000
Control Function Residual (Depth)	2.312	.431	.000	2.327	.431	.000
Occupation, MSA, Time, Firm Fixed Effects		YES			YES	
Controls		YES			YES	
Moderators		YES			YES	
Observations		1,493,148			1,493,148	
R ²		.271			.271	

Notes: Main variables are standardized; the changes in number of observations are due to missing values; we use 100 bootstrapping replications to calculate the SE

AI exposure and salary

Table 6 reports the salary results (full table shown in web appendix I2), main variables

⁴ Percentage effects are calculated as follows: (1) computed predict log(y) for each breadth-depth quadrant (e.g., high/high) using the estimated coefficient $s\beta_1$ (breadth), β_2 (depth), and β_3 (breadth × depth), at +/- SD for each variable; (2) took difference between the two quadrants; (3) convert difference to $\exp(\Delta) - 1$; (4) apply the percentage to the sample mean

are standardized and the highest VIF is 2.81 indicating multicollinearity is not a concern. In the main-effect model only (model 1), AI applicability breadth ($\beta = -2.449$, $p = .000$) and AI substitution depth ($\beta = -.902$, $p = .017$) are both negatively associated with posted salaries. In the interaction model, the negative interaction of breadth and depth indicates that the negative relationship between breadth and salary as both depth and breadth increase ($\beta = -.007$, $p = .000$). The salary results differ from the hiring results in an important respect. Breadth is positively associated with hiring but negatively associated with salary. This pattern seems consistent with firms expanding AI-enabled marketing work while reducing the wage premium associated with those roles, a dynamic that intensifies when AI can also perform a large share of the job's tasks. The negative interaction indicates that the salary penalty accompanying broader AI availability is larger for jobs where AI's substitution depth is also high. Broad AI availability can increase the ways firms conduct marketing, but it may also reduce scarcity rents as tasks become easier to execute, standardize, or reorganize. Depth is negative for both hiring and salary, consistent with the view that task-level substitutability places broader pressure on the labor-market value of affected roles. To translate our standardized estimates into economically interpretable terms, we estimate our model using $\log(\text{salary})$ as the dependent variable, while retaining standardized breadth and depth measures (results see column 8 of Table J2)⁴. The predicted salary gap between high-breadth/high-depth and low-breadth/low-depth is approximately 4.2%, roughly \$2,142 for a marketing position at the sample median salary (\$50,880).

Table 6. AI Depth, Breadth, and Salary

	Model 1			Model 2		
	Estimate	SE	P-value	Estimate	SE	P-value
Applicability Breadth	-2.429	.358	.000	-2.446	.373	.000
Substitution Depth	-.887	.384	.021	-.902	.377	.017
Breadth $\times$ Depth				-.007	.001	.000
Marketing Analytics $\times$ Breadth				-.064	.019	.001
Customer-management $\times$ Breadth				.077	.002	.000
AI $\times$ Breadth				.252	.014	.000

Sales × Breadth				.007	.002	.004
Brand Management × Breadth				.050	.014	.000
Marketing Operations × Breadth				.040	.005	.000
IMC × Breadth				.099	.004	.000
Marketing Analytics × Depth				-.022	.016	.179
Customer-management × Depth				.079	.004	.000
AI × Depth				-.163	.019	.000
Sales × Depth				.022	.002	.000
IMC × Depth				.055	.005	.000
Brand Management × Depth				-.075	.014	.000
Marketing Operations × Depth				-.044	.006	.000
Control Function Residual (Breadth)	2.395	.357	.000	2.400	.373	.000
Control Function Residual (Depth)	.875	.385	.023	.881	.377	.019
Occupation, MSA, Time, Firm		YES			YES	
Fixed Effects						
Controls		YES			YES	
Moderators		YES			YES	
Observations		2,297,723			2,297,723	
R ²		.698			.698	

Notes: Main variables are standardized; the changes in number of observations are due to missing values; we use 100 bootstrapping replications to calculate the SE.

Predictive Margins Across Breadth–Depth Profiles

The interaction results indicate that applicability breadth and substitution depth do not operate independently. To interpret the joint pattern more directly, we estimate predictive margins for four exposure profiles defined at one standard deviation above and below the mean of each dimension: low breadth–low depth, high breadth–low depth, low breadth–high depth, and high breadth–high depth in Table 7. This analysis allows us to distinguish whether AI exposure appears as labor-expanding augmentation, constrained substitution, or stronger wage and hiring pressure when both dimensions are high.

Table 7. Predictive Margins and Marginal Effects of AI Applicability Breadth and Substitution Depth

	Job Hiring	Posted Salary
Comparison	Estimate	Estimate
Panel A: Predictive Margins (at Mean ± 1SD)		
Low Breadth, High Depth vs Low Breadth, Low Depth	−3.835	−1.761
High Breadth, Low Depth vs Low Breadth, Low Depth	+3.692	−4.844
High Breadth, High Depth vs Low Breadth, Low Depth	−0.383	−6.633
High Breadth, Low Depth vs Low Breadth, High Depth	+7.527	−3.083

High Breadth, High Depth vs Low Breadth, High Depth	+3.452	−4.872
High Breadth, High Depth vs High Breadth, Low Depth	−4.075	−1.789

Panel B: Marginal Effects

Breadth ME at Low Depth (Mean−SD)	+1.846	−2.422
Breadth ME at High Depth (Mean+SD)	+1.726	−2.435
Depth ME at Low Breadth (Mean−SD)	−1.917	−0.881
Depth ME at High Breadth (Mean+SD)	−2.037	−0.894

The marginal-effects results qualify a simple automation-versus-augmentation interpretation. High breadth with low depth is associated with higher hiring but lower posted wages relative to low breadth with low depth. This pattern is consistent with labor-expanding but wage-depressing augmentation. Firms hire more when many AI technologies are applicable, but posted wages fall because some task content becomes less scarce. In contrast, high depth with low breadth is associated with lower hiring and lower wages relative to low breadth with low depth. This pattern is more consistent with constrained substitution: AI can perform a larger share of the job’s task bundle, exerting downward pressure on wages, but the limited breadth of applicable technologies may restrict AI adoption at scale.

The high-breadth, high-depth profile combines the labor expanding pressure associated with breadth and the substitution pressure associated with depth. Two comparisons clarify how these opposing forces interact. Relative to high breadth with low depth, high breadth with high depth is associated with lower hiring and lower wages, indicating that substitution depth weakens the hiring expansion otherwise associated with breadth. Relative to low breadth with high depth, however, high breadth with high depth is associated with higher hiring but lower wages, suggesting that breadth still supports some scale expansion even when depth is high. Thus, the negative interaction does not imply uniform employment contraction across all high-exposure jobs. Instead, joint exposure most clearly appears as wage compression and as an attenuation of the positive hiring effect associated with breadth.

The comparison between low breadth–low depth and high breadth–high depth further reinforces this interpretation. These two profiles are not clearly different in hiring, even though high breadth–high depth carries a larger wage penalty. The strongest evidence for joint AI exposure is therefore not that the most AI-amenable jobs experience the largest absolute decline in hiring. Rather, the evidence indicates that breadth and depth create offsetting hiring forces while jointly intensifying wage pressure. This pattern underscores the value of a two-dimensional measure of AI exposure.

Post-hoc Analysis of Marketing Skill Moderation

We conducted a post hoc analysis of patterns of heterogeneity. The skill-interaction estimates help explain why the main effects of AI exposure differ across marketing roles (all estimations shown in Tables 5 and 6). They show that AI exposure does not simply raise or lower the value of marketing human capital. It reprices different skills depending on whether AI expands the technology set available to marketing labor or reaches the important work that supports the role.

Three regularities stand out. First, the same skill can have different salary implications depending on AI applicability breadth and AI substitution depth. AI skills provide the clearest example. In the salary model, AI skills attenuate the negative effect of breadth ($\beta = .252$, $p = .000$) but amplify the negative effect of depth ($\beta = -.163$, $p = .000$), suggesting that AI fluency appears valuable when workers use a broader set of AI technologies but offers less protection when AI can perform the core task bundle itself. Brand Management and Marketing Operations show similar wage patterns. Brand Management positively moderates breadth ($\beta = .050$, $p = .000$), but negatively moderates depth ($\beta = -.075$, $p = .000$). Marketing Operations also positively moderates breadth ($\beta = .040$, $p = .000$), but negatively moderates depth ($\beta = -.044$, p

= .000). These findings suggest that skills can retain value when AI broadens what workers can do, while losing wage protection when AI can perform the underlying work.

Second, salary protection and hiring protection are not the same. Customer management, sales, and integrated marketing communication skills are more clearly wage-protective than hiring-protective. Customer-management skills positively moderate both breadth ($\beta = .077$, $p = .000$), and depth ($\beta = .079$, $p = .000$), in the salary model, but do not significantly moderate either dimension in the hiring model. Sales skills also attenuate salary penalties under breadth ($\beta = .007$, $p = .000$), and depth ($\beta = .022$, $p = .000$). Yet in hiring, sales skills show a symmetric pattern, in which it positively moderates breadth ($\beta = .012$, $p = .011$), and negatively moderates depth ($\beta = -.056$, $p = .000$). Integrated marketing communication skills are wage-protective under both breadth ($\beta = .099$, $p = .000$), and depth ($\beta = .055$, $p = .000$), but negatively moderate hiring under breadth ($\beta = -.067$, $p = .000$), and depth ($\beta = -.043$, $p = .000$). Overall, firms may continue to pay for relational, persuasive, and communication-oriented skills in the roles they retain, even if AI reduces the number of jobs posted for some of these roles, particularly under higher substitution depth pressure.

Third, the Marketing Operations moderation finding shows that AI's labor-market effects do not follow a simple routine-versus-non-routine distinction. It is the most consistent hiring-protective category, positively moderating both breadth ($\beta = .100$, $p = .000$) and depth ($\beta = .027$, $p = .006$) in the hiring model. This is notable because marketing operations encompass structured, codifiable activities such as CRM, campaign automation, measurement routines, and performance reporting. If codifiability alone drove substitution, this category should face stronger hiring pressure. The results suggest instead that AI changes the composition of marketing operations work. Firms may automate parts of the task bundle while increasing

demand for workers who integrate AI into campaign systems, data flows, dashboards, governance routines, and cross-functional workflows. Marketing operations is therefore hiring-protective, not because they are shielded from AI, but because they provide the infrastructure through which AI-enabled marketing is implemented.

The salary results are consistent with this interpretation. Marketing Operations attenuates salary penalties under breadth ($\beta = .040$, $p = .000$) but amplifies them under depth ($\beta = -.044$, $p = .000$). Thus, firms may need more marketing operations capacity as AI spreads, even as AI weakens the wage premium for codifiable operational tasks.

Brand Management has a more conditional pattern. It positively moderates breadth in hiring ($\beta = .098$, $p = .000$), and salary ($\beta = .050$, $p = .000$). It negatively moderates depth in hiring ($\beta = -.054$, $p = .021$), and salary ($\beta = -.075$, $p = .000$). Brand-related skills appear valuable when AI expands the technologies available for market analysis, content development, coordination, and brand stewardship. They are less protected when AI reaches deeper into the task bundle. This does not imply that brand judgment is obsolete. It suggests that the labor-market value of brand skills depends on whether the role centers on market interpretation and responsibility for brand direction or on executional tasks that AI can produce at scale.

The final pattern is that technical skills do not automatically shield workers from exposure to AI. Marketing analytics skills are negatively moderate breadth in both hiring ($\beta = -.091$, $p = .001$) and salary ($\beta = -.064$, $p = .001$). Their interactions in depth are not significant at the 10% level in the full models. Analytics remains central to marketing, but many analytics tasks use structured data, measurable feedback, and repeatable decision rules. Those same features make parts of analytics work amenable to AI. AI skills show related caution. They protect wages under breadth but not depth, and they do not significantly protect hiring under

depth. Technical fluency is useful when it complements judgment. It is less protective when the work itself becomes machine-executable.

Taken together, the moderator results call into question any simple ranking of “safe” and “exposed” marketing skills. The value of a skill depends on the AI channel, the task bundle, and the outcome being examined. Salary protection is strong where work relies on relational judgment, persuasion, communication, and contextual interpretation. Hiring protection is strongest where skills help firms operationalize AI-enabled marketing work. Therefore, the post hoc results point to a more precise view of marketing human capital. AI does not merely reduce or expand demand for marketing labor. It reallocates demand across task bundles and reprices the skills attached to those bundles.

Robustness Checks

We conducted several additional robustness analyses to check the robustness of the results, and the full results are shown in web appendix Tables J1 and J2. Across all specifications, the main findings remain consistent. First, to rule out concerns that the results are driven by control variables or by the control-function correction itself, we estimate models that exclude control variables and control-function residuals. As shown in column 1 of Table J1 and J2, our results are robust to a simpler model. Second, O*NET occupation-level importance scores, which we adopt to maintain consistency with prior work using measures of AI exposure, may not capture the variations across postings within the same occupation and are updated infrequently, potentially failing to capture shifts in task importance during the sample period. We address these concerns with an alternative operationalization, task position within each posting, where earlier-listed tasks are assumed to carry greater importance (i.e., position is inversely related to importance). Because this measure is computed directly from each posting rather than

assigned from a fixed external source, it is posting-specific and time-varying, flexing both across postings within an occupation and across our sample period. Column 2 of Tables J1 and J2 shows that our independent effects of breadth and depth are robust to the alternative measurement of importance score. Third, we replace the baseline fixed-effects structure with more demanding interacted fixed effects, including MSA-by-occupation and firm-by-time fixed effects. As shown in Tables J1 and J2 columns 3 and 4, our results remain consistent with the inclusion of higher-order interacted fixed effects. Fourth, to further examine whether our AI exposure measure retains independent associations with outcomes after accounting for the previous measure and captures information beyond the existing GPT-exposure measure. To do so, we include the exposure score from Eloundou et al. (2024), a well cited measure published in *Science*, as an additional regressor into our model. As shown in column 5 of Tables J1 and J2, the results remain consistent with the main analyses and adding breadth, depth, and their interaction to the model that already includes the Eloundou et al (2024) measure improves the model fit ($F\text{-statistics} = 1480.497$, $p = .000$). These results provide evidence that the two-dimensional decomposition captures extra information not accounted for by the established measure. Fifth, to address concerns about selection into salary disclosure, we re-estimate our main models separately for postings from states with mandatory disclosure laws and those without. Results in column 6 and 7 of Table J2 show that our results are robust across different disclosure samples. Finally, because our main job hiring analysis relies on postings that report salary information, the results could be sensitive to this sample restriction. Therefore, we estimate the model using the full posting sample rather than the salary-only sample. Relatedly, our results may also depend on the window used to aggregate hiring activity. To address this, we aggregate hiring at the monthly rather than quarterly level. As shown in Table J1, columns 6, 7

and 8, the results are robust to alternative sample and aggregate windows.

Discussion

Concurrent marketing-employment studies identify effects from a single ChatGPT shock and treat AI exposure as essentially one-dimensional (Luo, Miao, and Sudhir 2026; Zhu et al. 2026). Our results suggest that this approach masks meaningful heterogeneity in how AI relates to marketing labor-market outcomes. We contribute by decomposing AI exposure into two measured dimensions. A marketing job may be exposed because many AI technologies can be applied to its tasks, which we define as applicability breadth, or because AI can perform a large share of the job's core task bundle, which we define as substitution depth. By linking the task content of marketing job postings to the functional content of AI patents, we examine how AI maps onto marketing work at the task level, how applicability breadth and substitution depth relate to wages and hiring, and how marketing-specialized skills condition these relationships.

The results show that breadth and depth are related but distinct dimensions of Marketing AI Exposure. They vary across jobs and industries, generate meaningful off-diagonal cases, and have different associations with labor-market outcomes. Breadth is positively associated with hiring but negatively associated with salary. Depth is negatively associated with both hiring and salary. The negative interaction between breadth and depth further indicates that the two dimensions do not operate independently. However, the marginal analysis qualifies a simple displacement interpretation. High-breadth, high-depth jobs do not experience uniformly lower hiring than low-breadth, low-depth jobs. Instead, joint exposure is most clearly associated with wage compression and with a weakening of the hiring expansion otherwise associated with breadth. Interpreted through this lens, breadth captures the feasibility margin of AI adoption

because more technologies are available for a job's task bundle, whereas depth captures the substitution-value margin because AI can perform more of the work itself.

Breadth associated with higher hiring but lower wages exemplifies labor-expanding augmentation, in which broad AI applicability may reduce the cost or risk of executing marketing activities that support expanded output, workflow redesign, or creation of adjacent tasks, even as posted wages decline. This interpretation is consistent with evidence that generative AI reorganizes rather than eliminates marketing work, expanding orchestration-oriented functions while contracting production-oriented ones (Luo, Miao, and Sudhir 2026). Depth's association with lower hiring and lower wages, by contrast, reflects substitution rather than augmentation, consistent with substitution-value margin.

Theoretical Implications

Concurrent studies identify marketing-employment effects from the emergence of ChatGPT while treating AI exposure as largely one-dimensional (Luo, Miao, and Sudhir 2026; Zhu et al. 2026). Our findings show why this approach obscures theoretically important heterogeneity. By distinguishing AI applicability breadth from AI substitution depth, and by linking marketing-job tasks to AI-patent functions, we separate two mechanisms that are often conflated: the availability of AI technologies relevant to a job and the extent to which those technologies can perform the job's core task content. Breadth and depth are related, but they are not interchangeable. They produce meaningful off-diagonal cases and have different relationships with wages and hiring. This distinction helps explain why AI exposure does not imply a uniform pattern of displacement, augmentation, or wage growth.

The findings refine how AI exposure should be theorized in marketing labor markets. Breadth reflects the feasibility of AI adoption, because a wider set of applicable technologies

lowers the cost and uncertainty of redesigning marketing activity. Depth reflects substitution pressure, because it captures the extent to which AI can perform the core task content of a job. This distinction explains why augmentation need not be wage-enhancing. Breadth is associated with higher hiring but lower wages, suggesting that AI availability may expand marketing activity or shift work toward orchestration while reducing the scarcity value of some human inputs. Depth, by contrast, is associated with lower hiring and lower wages, consistent with stronger substitution pressure. The negative interaction further shows that these mechanisms are interdependent. When AI is both widely applicable and capable of performing core tasks, the hiring expansion associated with breadth weakens and wage compression intensifies. Thus, AI's employment consequences depend not simply on whether a job is exposed to AI, but on whether exposure reflects technological availability, task performability, or their combination.

The marginal analysis further qualifies a simple displacement account. High-breadth, high-depth jobs do not exhibit uniformly lower hiring than low-breadth, low-depth jobs. Instead, joint exposure is more clearly associated with wage compression and with a weakening of the positive hiring relationship observed for breadth. This pattern suggests that AI may reorganize marketing work rather than eliminate it. The result is consistent with evidence that generative AI expands orchestration-oriented marketing functions while contracting more production-oriented ones (Luo, Miao, and Sudhir 2026). More broadly, the findings distinguish labor-expanding augmentation from wage-increasing augmentation. AI can increase demand for some forms of marketing work while simultaneously reducing the wage premium attached to tasks that become more standardized, more executable by technology, or less dependent on scarce human expertise.

Finally, the analysis of marketing-specialized skills extends this theoretical account by showing that AI's labor-market effects are conditioned by the marketing human capital

embedded in a job. Skills do not merely indicate whether a job is more or less exposed to AI. They shape how breadth and depth translate into wages and hiring. This implies that the consequences of AI depend on the interaction between technological capability and the organization of marketing work. Jobs that rely on coordination, integration, judgment, and orchestration may remain hiring-relevant even as wage pressure rises in tasks that AI can increasingly perform. Theoretical accounts of AI and marketing work therefore need to move beyond aggregate exposure and examine how specific capabilities interact with the task structure and skill content of jobs.

Managerial Implications

Firms should not evaluate AI exposure at the level of broad job titles alone. Managers need to assess two distinct questions. How many AI technologies are applicable to a role? How deeply do those technologies reach into the role's core responsibilities? A high-breadth, low-depth role may be a candidate for AI-enabled workflow redesign. In such roles, AI can support more campaigns, analyses, experiments, content variants, and customer interactions, even if it weakens the wage premium attached to some executional tasks.

Roles with high substitution depth require a different response. When AI reaches core task content, firms should examine which responsibilities remain distinctly human and redesign jobs around judgment, coordination, interpretation, and accountability. The marginal analysis suggests that high-breadth, high-depth roles do not necessarily experience the largest hiring decline, but they do face the strongest wage pressure. Managers should therefore avoid equating AI exposure only with headcount reduction. The more immediate issue may be job redesign, changing skill requirements, and compression in the wage value of tasks that AI can perform.

The findings also suggest that managers should distinguish between roles needed to

deploy AI and roles whose tasks AI can perform. Marketing operations is the clearest example of a hiring-protective category because firms still need employees to integrate AI into campaign systems, dashboards, governance routines, and cross-functional workflows. Brand management and AI-related skills may protect wages when AI broadens what workers can do, but they become less protective when AI reaches the core task bundle. Skill development should therefore move beyond AI tool training alone. Firms should pair AI fluency with market judgment, customer understanding, persuasive communication, cross-functional coordination, and responsibility for decisions.

Implications for Marketing Education

Marketing curricula should combine AI literacy with skills that complement AI-enabled work. Training in AI tools, analytics, automation, and experimentation remains important, but our findings indicate that technical fluency may not preserve labor-market value when AI can perform core job tasks. The moderator results suggest greater emphasis on skills involving customer interaction, persuasion, market interpretation, and coordination.

These results also show that skills can relate differently to wages and hiring. Relational and customer-facing skills more often attenuate negative wage associations, whereas operational and coordination skills more often attenuate negative hiring associations. Marketing programs should therefore avoid treating all specialized skills as equally valuable across labor-market outcomes. This implication is particularly relevant given evidence that exposure to generative AI may reduce demand for some master's-trained workers in business and other non-STEM occupations (Cortes, Dellarocas, and Wang 2026).

Limitations and Future Research

The study has limitations that create opportunities for future research. First, our measures

are based on job postings and patents. Job postings reflect employer demand and advertised requirements, but they do not capture all tasks performed after hiring or all informal changes in work organization. Patents provide a systematic record of AI capabilities, but not all patented technologies are commercialized or adopted. Future research could combine postings with firm-level AI adoption data, internal workflow data, or employee-level career histories.

Second, the empirical design uses rich fixed effects and a control-function approach, but our identification strategy does not exploit random variation in AI exposure created by a randomized controlled trial or by a natural experiment. Future work could exploit shocks to AI availability or adoption within firms. Finally, the moderator analysis is post hoc. It identifies heterogeneity in skill pricing, hiring and wage protection, but it does not establish the mechanism through which each skill category changes worker productivity or firm demand. Relatedly, our post-hoc analysis treats skill categories as exogenous and we view results as associative and future work to address causal identification.

Reference

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022), “Artificial Intelligence and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 40 (S1), S293–340.
- Acemoglu, Daron and Pascual Restrepo (2019), “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33 (2), 3–30.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb (2018), “Human Judgment and AI Pricing,” *AEA Papers and Proceedings*, 108, 58–63.
- Ahearne, Michael, Eli Jones, Adam Rapp, and John Mathieu (2008), “High Touch Through High Tech: The Impact of Salesperson Technology Usage on Sales Performance via Mediating Mechanisms,” *Management Science*, 54 (4), 671–85.
- Anders, George (2026), “Is an AI Jobs Apocalypse Coming? Three Economists Square Off,” *Wall Street Journal* (June 2026), <https://www.wsj.com/economy/jobs/ai-jobs-economists-f787105d>.
- Arnold, David, Simon Quach, and Bledi Taska (2025), “The Impact of Pay Transparency in Job Postings on the Labor Market,” National Bureau of Economic Research.

- Arntz, Melanie, Terry Gregory, and Ulrich Zierahn (2017), “Revisiting the risk of automation,” *Economics Letters*, 159, 157–60.
- Autor, David H. and David Dorn (2013), “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103 (5), 1553–97.
- Autor, David H., Frank Levy, and Richard J. Murnane (2003), “The Skill Content of Recent Technological Change: An Empirical Exploration*,” *The Quarterly Journal of Economics*, 118 (4), 1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson (2024), “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 151, 103745.
- Bansal, Naresh, Kissan Joseph, Minghui Ma, and M. Babajide Wintoki (2017), “Do CMO Incentives Matter? An Empirical Investigation of CMO Compensation and Its Impact on Firm Performance,” *Management Science*, 63 (6), 1993–2015.
- Berger, Jonah, Ashlee Humphreys, Stephan Ludwig, Wendy W. Moe, Oded Netzer, and David A. Schweidel (2020), “Uniting the Tribes: Using Text for Marketing Insight,” *Journal of Marketing*, 84 (1), 1–25.
- Bindley, Katherine (2026), “Big Tech Has Suddenly Flipped on the AI Jobs Wipeout Scenario,” *Wall Street Journal*, (July 2026), <https://www.wsj.com/tech/ai/ai-workers-tech-ceos-job-losses-afc71e15>.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025), “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.”
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025), “Generative AI at work,” *The Quarterly Journal of Economics*, 140 (2), 889–942.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018), “What Can Machines Learn and What Does It Mean for Occupations and the Economy?,” *AEA Papers and Proceedings*, 108, 43–47.
- Burke, Raymond R, Maximilian Matthe, and Alex Leykin (2025), “Assessing Artificial Marketing Intelligence,” MSI working paper series.
- Cortes, Patricia, Chrysanthos Dellarocas, and Qi Wang (2026), “Generative AI and the Informational Value of Educational Credentials: Evidence from the Master’s Margin,” SSRN Scholarly Paper, Rochester, NY: Social Science Research Network.
- Cutter, Chip and Haley Zimmerman (2025), “CEOs Start Saying the Quiet Part Out Loud: AI Will Wipe Out Jobs,” *Wall Street Journal*, (July 2025), <https://www.wsj.com/tech/ai/ai-white-collar-job-loss-b9856259>.
- Damioli, Giacomo, Vincent Van Roy, Daniel Vertesy, and Marco Vivarelli (2023), “AI technologies and employment: micro evidence from the supply side,” *Applied Economics Letters*, 30 (6), 816–21.
- Davenport, Thomas, Abhijit Guha, Dhruv Grewal, and Timna Bressgott (2020), “How artificial intelligence will change the future of marketing,” *Journal of the Academy of Marketing Science*, 48 (1), 24–42.
- Dell’Acqua, Fabrizio, Edward McFowland, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Candelon, and Karim 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), 403–23.
- Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska (2023), “Generative AI and Firm Values,” SSRN Scholarly Paper, Rochester, NY: Social Science Research Network.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024), “GPTs are GPTs: Labor market impact potential of LLMs,” *Science*, 384 (6702), 1306–8.
- Felten, Edward, Manav Raj, and Robert Seamans (2021), “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses,” *Strategic Management Journal*, 42 (12), 2195–2217.
- Fleming, Lee and Olav Sorenson (2004), “Science as a map in technological search,” *Strategic Management Journal*, 25 (8–9), 909–28.
- Furneaux, Brent (2012), “Task-Technology Fit Theory: A Survey and Synopsis of the Literature,” in *Information Systems Theory: Explaining and Predicting Our Digital Society, Vol. 1*, Y. K. Dwivedi, M. R. Wade, and S. L. Schneberger, eds., New York, NY: Springer, 87–106.
- Garvey, Aaron M., TaeWoo Kim, and Adam Duhachek (2023), “Bad News? Send an AI. Good News? Send a Human,” *Journal of Marketing*, 87 (1), 10–25.
- Germann, Frank, Peter Ebbes, and Rajdeep Grewal (2015), “The Chief Marketing Officer Matters!,” *Journal of Marketing*, 79 (3), 1–22.
- Goodhue, Dale L. and Ronald L. Thompson (1995), “Task-Technology Fit and Individual Performance,” *MIS Quarterly*, 19 (2), 213–36.
- Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller (2025), “Artificial Intelligence and the Labor Market,” Working Paper Series, Working Paper, National Bureau of Economic Research.
- Hartmann, Jochen, Yannick Exner, and Samuel Domdey (2025), “The power of generative marketing: Can generative AI create superhuman visual marketing content?,” *International Journal of Research in Marketing*, 42 (1), 13–31.
- Howard, Matt and Joseph Hair (2023), “Integrating the Expanded Task-technology Fit Theory and the Technology Acceptance Model: A Multi-wave Empirical Analysis,” *AIS Transactions on Human-Computer Interaction*, 15 (1), 83–110.
- Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024), “The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market,” *Organization Science*.
- Humphreys, Ashlee and Rebecca Jen-Hui Wang (2018), “Automated Text Analysis for Consumer Research,” *Journal of Consumer Research*, 44 (6), 1274–1306.
- Kaiser, Maximilian and Christian Schulze (2026), “Frontiers: ChatGPT Referrals to E-Commerce Websites: How Do LLMs Compare Against Traditional Channels?,” *Marketing Science*.

- Kapoor, Anuj and Madhav Kumar (2025), “Frontiers: Generative AI and Personalized Video Advertisements,” *Marketing Science*, 44 (4), 733–47.
- Keane, Michael and Timothy Neal (2023), “Instrument strength in IV estimation and inference: A guide to theory and practice,” *Journal of Econometrics*, 235 (2), 1625–53.
- Kumar, V., Abdul R. Ashraf, and Waqar Nadeem (2024), “AI-powered marketing: What, where, and how?,” *International Journal of Information Management*, 77, 102783.
- Lawrence, Justin M., Lisa K. Scheer, Andrew T. Crecelius, and Son K. Lam (2021), “Salesperson Dual Agency in Price Negotiations,” *Journal of Marketing*, 85 (2), 89–109.
- Lemmens, Aurélie, Jason M. T. Roos, Sebastian Gabel, Eva Ascarza, Hernán A. Bruno, Brett R. Gordon, Ayelet Israeli, Elea McDonnell Feit, Carl F. Mela, and Oded Netzer (2025), “Personalization and targeting: how to experiment, learn & optimize,” *International Journal of Research in Marketing*.
- Li, Bingqing, Edward Yuhang Lai, and Xin (Shane) Wang (2025), “From Tools to Agents: Meta-Analytic Insights into Human Acceptance of AI,” *Journal of Marketing*.
- Longoni, Chiara and Luca Cian (2022), “Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The ‘Word-of-Machine’ Effect,” *Journal of Marketing*, 86 (1), 91–108.
- Lu, Zipei and P.K. Kannan (2025), “AI for Customer Journeys: A Transformer Approach,” *Journal of Marketing Research*, 00222437251347268.
- Luo, Xueming, Jason Miao, and K Sudhir (2026), “The Production-to-Orchestration Shift: How Generative AI Reorganizes Knowledge Work,” SSRN Scholarly Paper, Rochester, NY: Social Science Research Network.
- Mahmood, Mo Adam, Laura Hall, and Daniel Leonard Swanberg (2001), “Factors Affecting Information Technology Usage: A Meta-Analysis of the Empirical Literature,” *Journal of Organizational Computing and Electronic Commerce*, 11 (2), 107–30.
- Mu, Jifeng and Jonathan Z. Zhang (2025), “Artificial intelligence marketing usage and firm performance,” *Journal of the Academy of Marketing Science*.
- Nath, Shanjukta, Jinsoo Yeo, Anandhi Bharadwaj, and Sundar G. Bharadwaj (2026), “Pricing of Multidimensional Skills in IT Labor Market: Complementarities between AI and Traditional IT Skills,” SSRN Scholarly Paper, Rochester, NY: Social Science Research Network.
- Netzer, Oded, Ronen Feldman, Jacob Goldenberg, and Moshe Fresko (2012), “Mine Your Own Business: Market-Structure Surveillance Through Text Mining,” *Marketing Science*, 31 (3), 521–43.
- Seetharaman, Deepa and Georgia Wells (2023), “Tech Leaders Say AI Will Change What It Means to Have a Job,” *Wall Street Journal* (October 2023), <https://www.wsj.com/tech/ai/tech-leaders-say-ai-will-change-what-it-means-to-have-a-job-2dd556fb>.
- Singh, Siddharth Shekhar, Ravi Sen, and Sharad Borle (2022), “Online Training of Salespeople: Impact, Heterogeneity, and Spillover Effects,” *Journal of Marketing Research*, 59 (1), 230–49.

- Srivastava, Chandra, Saim Kashmiri, and Vijay Mahajan (2023), “Customer Orientation and Financial Performance: Women in Top Management Teams Matter!,” *Journal of Marketing*, 87 (2), 190–209.
- Sutton, Richard S. and Andrew G. Barto (2018), *Reinforcement learning: An introduction*, 2nd ed, Reinforcement learning: An introduction, 2nd ed, Cambridge, MA, US: The MIT Press, xxii, 526.
- Tully, Stephanie M., Chiara Longoni, and Gil Appel (2025), “Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity,” *Journal of Marketing*, 89 (5), 1–20.
- Wang, Mengxin, Dennis J. Zhang, and Heng Zhang (2026), “Large Language Models for Market Research: A Data-Augmentation Approach,” *Marketing Science*.
- Webb, Michael (2019), “The Impact of Artificial Intelligence on the Labor Market,” SSRN Scholarly Paper, Rochester, NY.
- Wedel, Michel and P.K. Kannan (2016), “Marketing Analytics for Data-Rich Environments,” *Journal of Marketing*, 80 (6), 97–121.
- Whitler, Kimberly A., Ryan Krause, and Donald R. Lehmann (2018), “When and how Board Members with Marketing Experience Facilitate Firm Growth,” *Journal of Marketing*, 82 (5), 86–105.
- Wooldridge, Jeffrey M. (2015), “Control Function Methods in Applied Econometrics,” *The Journal of Human Resources*, 50 (2), 420–45.
- Zhu, Ruizhi, Samsun Knight, Anatoli Colicev, and Yakov Bart (2026), “The Effect of AI on Marketing Employment: Evidence from 110m+ Online Employment Records,” SSRN Scholarly Paper, Rochester, NY: Social Science Research Network.

Marketing AI Exposure: Applicability Breadth, Substitution Depth, and Job Outcomes

Web Appendix

Web Appendix A: Marketing Occupation Groups Selection	2
Web Appendix B: Dependency Parsing and Verb-Noun Pair Extraction	4
Web Appendix C: Importance Score from O*NET Database	7
Web Appendix D: Real example of Pipeline.....	8
Web Appendix E: Selection of Hierarchy Depth and Measurement Validation	9
Web Appendix F: Descriptive Statistics by Occupation Groups and Industries	14
Web Appendix G: Skill Classification.....	16
Web Appendix H: Control Function First-Stage Results	17
Web Appendix I: Full Results.....	19
Web Appendix J: Robustness Check	23

These materials have been supplied by the authors to aid in the understanding of their paper. The AMA is sharing these materials at the request of the authors.

Web Appendix A: Marketing Occupation Groups Selection

The Lightcast database encompasses job postings across more than 180 occupation groups and 800 distinct sub-occupations. These occupation groups are not pre-classified as marketing-related or non-marketing related, necessitating a systematic procedure to identify those relevant to our study. To identify marketing-related occupation groups, we employed a two-step procedure. In the first step, authors independently reviewed the full list of 180 different occupation groups and reached a consensus to exclude those occupation groups that are unambiguously unrelated to marketing (e.g., Physicians and Loan Officers). This step reduced the list to 32 occupation groups, which were then submitted to expert evaluation in the second step. In the second step, we conducted a survey of 6 marketing professionals recruited based on their expertise in marketing practice and research, each of whom independently reviewed the description of each remaining occupation group and indicated whether it qualified as a marketing-related role (see Table A1 for respondent profiles).

Table A1. Respondent Profiles

Participant	Position	Years in Experiences
1	Marketing Strategy Manager	6 – 10 years
2	Marketing Executive	More than 10 years
3	Marketing Professor	More than 10 years
4	Senior Project Manager	5 years
5	Sales Manager	5 years
6	Entrepreneur	More than 10 years

An occupation group was retained in the final sample only if more than 50% of respondents classified it as marketing-related. In cases where exactly 50% of respondents agreed, the authors collectively reviewed the occupation group and discussed the occupation group to reach a consensus decision. This procedure yielded a final set of 24 marketing-related occupation groups (see Table A2 for the full list of evaluated groups and their classification outcomes).

Table A2. Occupation Groups and Results

Occupation Groups	Marketing-related
Agriculture Managers	×
Business Analysis	√
Fundraising and Grants	√
Customer Service Managers	√
Customer Service Representatives	√
Product Design	×
Writing and Language	Only copywriter
Financial Sales	√
Insurance Claims and Sales	√
Health Education and Counseling	×
Business Intelligence	√
Front-End Application Design	×
Communication and Public Relations	√
Marketing Managers	√
Marketing Specialists	√
Account Managers and Executives	√
Energy Sales	√
Merchandising	√
Non-Technical Sales	√
Promotional Modeling	√
Retail Managers	√
Sales Managers	√
Retail Sales	√
Technical Sales	√
Billing and Collections	×
Visual Design	√
Hospitality	√
Hospitality Managers	√
Recreation and Gaming	×
Data Analysis and Mathematics	√
General Research	×
Social Analysis	×
Delivery Driver and Messengers	×

Web Appendix B: Dependency Parsing and Verb-Noun Pair Extraction

Dependency parsing is a syntactic analysis technique in Natural Language Processing (NLP) that represents each sentence as a directed dependency tree, in which each word is linked by labeled arcs that encode their grammatical relationships, with the main verb serving as the root of the tree. Unlike sequential or constituency-based representations, this structure aligns more directly with the semantic content of a sentence, making it well suited for extracting meaningful action-object relationships. This approach is particularly well-suited for information extraction tasks, as it captures the core semantic relationships in a sentence regardless of surface word order or syntactic complexity (Martin and Kushwaha 2024). In the dependency tree, each word is represented as a node, and each directed edge denotes both the type of grammatical relationship and the direction of dependency between two words.

As dependency parsing is particularly useful in information extraction applications, in this study, we leverage it to extract verb-noun pairs from marketing job responsibility descriptions and AI patent abstracts, thereby identifying the core actions embedded in each source. Specifically, we use the SpaCy package for Python (Montani et al. 2023), which achieves 91.85% accuracy on the dependency parsing benchmark in NLP literature, making it a reliable tool for our analysis (Webb 2019). Beyond identifying head-dependent relationships, the algorithm also classifies the grammatical relations each dependent holds with respect to its head (e.g., subject, direct object) and assigns the grammatical type (e.g., noun, verb) of each word by using a part of speech (POS) library, both of which are used in our verb-noun pair extraction procedure described below.


To illustrate how we extract verb-noun pairs from sentences via dependency parsing, we walk through the following example sentence drawn from actual job posting: “*Create*

comprehensive multimedia communication plans.” First, we tagged each word with its POS and dependency relation and lemmatized each word, as shown in Table B1. The ROOT label identifies the main predicate of the sentence, which has no governing head word of its own. Each remaining word is linked to a head word, with the direction and type of each dependency relation recorded in the Head and Dependency Relation columns of Table B1. Figure B1 visualizes the resulting dependency tree, in which each arrow points from the head word to its dependent and the label on each arrow indicates grammatical relation type.

Table B1. Dependency Parsing Example

Word	Lemmatized Word	POS	Dependency Relation	Head
Create	create	VERB	ROOT	ROOT
comprehensive	comprehensive	ADJ	amod	plans
multimedia	multimedia	NOUN	compound	plans
communication	communication	NOUN	compound	plans
plans	plan	NOUN	dobj	create

Figure B1. Dependency Parsing Graph Example



Building on this parsed representation, we extract verb-noun pairs from the dependency tree through the following three steps (Webb 2019). First, we remove stop words and auxiliary verbs that carry little semantic content and economic meaning (e.g., “have,” “be”). Second, we identify each remaining verb from the sentence and extract its direct object as determined by the dependency parser, storing the resulting verb-noun pair. For example, from the example sentence

above, we extract “create + plan.” If no direct object is present, we instead extract the prepositional object as a substitute (e.g., extracting “engage + audience” from “engage with audience”). Third, each extracted noun is mapped to a generalized semantic category using the WordNet hypernym hierarchy (Miller 1995). Specifically, we traverse each noun upward along its hypernym path and replace it by the concept at a fixed depth of level 6 from the WordNet root (see the validation of selection in the web appendix E). This normalization step maps semantically related but lexically distinct nouns into a shared abstract category, enabling cross-source comparison at an appropriate level of generality and allowing us to identify the high-level general functions of AI technology and goals of different tasks. For example, both “communicate + users” and “communicate + clients” would be part of the verb-noun pair “communicate + customer,” allowing them to be treated as equivalent across job postings and patent abstracts.

Web Appendix C: Importance Score from O*NET Database

Because marketing occupations vary in the relative importance of their constituent tasks, we weight each task by its importance to the occupation when aggregating depth and breadth scores. We obtain the task importance score from O*NET, which assigns each task a score ranging from 1 to 5 to indicate its centrality to the occupation. However, tasks extracted from job postings are described in varied linguistic forms that do not always correspond exactly to O*NET task descriptions. We therefore link each task in our database to its most similar O*NET task using cosine similarity.

To obtain an important score of each task in our database, we link each task to its closest counterpart in the O*NET task database using cosine similarity. Given that Lightcast provides O*NET occupation codes for each job posting, we first link each job posting to its corresponding O*NET occupation and restrict similarity matching to tasks within that occupation, thereby avoiding spurious across-occupation matches and improving assignment accuracy. We then employ the GTE-large embedding model (developed by the Alibaba DAMO Academy), which has demonstrated superior performance on semantic similarity benchmarks compared with other models of comparable scale (Hampole et al. 2025). This model encodes each task description into a 1024-dimensional vector, enabling fine-grained semantic comparison between job posting tasks and O*NET task descriptions. For each task in our database, we calculate its cosine similarity with all O*NET tasks within the matched occupation and assign the importance score of the best-matching O*NET task with the highest similarity score.

Web Appendix D: Real example of Pipeline

Example job responsibilities of sales representative position

Primary responsibilities of an outside sales representative include:

- Driving Sales Growth of our automotive products within a designated sales territory
- Educating current and potential customers on the features and benefits of our products.
- Maintaining customer relationships.

Step 1: Extract verb-noun pair from each task

1. Driving Sales Growth of our automotive products within a designated sales territory → (drive, growth)
2. Educating current and potential customers on the features and benefits of our products → (educate, customers)
3. Maintaining customer relationships → (maintain, relationship)

Step 2: Calculate the applicability breadth and substitution depth at verb-noun pair level

Pair	Applicability Breadth	Substitution Depth
drive, growth	0	0
educate, customer	4	1
maintain, relationship	244	1

Step 3: Assign an importance score to each task

Task	O*NET Task (same occupation code)	Importance Score
1	Initiate sales campaigns to meet sales and production expectations.	3.7
2	Advise customers on product usage to improve production.	3.65
3	Identify prospective customers, using business directories, leads from existing clients, participation in organizations, or trade show or conference attendance.	3.79

Step 4: Calculate the applicability breadth and substitution depth at job-level

Publish date: June 2021; Total number of AI patents before June 2021 = 1,077,375
Applicability breadth = 0.0000807; Substitution depth = 0.691

Web Appendix E: Selection of Hierarchy Depth and Measurement Validation

Part 1. Selection of Hierarchy Depth and Verb-Noun Validation

First, we conducted 200 participants recruited from Prolific, an online crowdsourcing platform, to validate the selection of hierarchy depth. This study aimed to select the optimal hierarchy depth and validate the verb-noun method. Participants evaluated the representativeness of verbs, nouns, and verb-noun pairs extracted at different hierarchy levels. We randomly selected 20 sentences from the dataset and extracted verbs and nouns at WordNet levels 3 through 6. Each participant was randomly assigned to one of four levels and rated the verb-noun pair on three dimensions using seven-point Likert scales (see Table E1 for an example of what the participant saw at each level).

Table E1. Example of Verb-Noun Validation

	Level 1	Level 2	Level 3	Level 4
Task Statement	Task Statement: Actively and consistently supports all efforts			
Verb-Noun	Support + Event	Support + Act	Support + Activity	Support + Work
Questions	Please rate the accuracy of the verb "[Verb]" in describing the action performed in this task: (1) Extremely inaccurate — (4) Neutral — (7) Extremely accurate To what degree is the noun "[Noun]" representative of the primary object or outcome involved in this task? (1) Not at all — (4) Moderate — (7) Highly Representative To what extent does the pair "[verb noun]" capture the nature of the task? (1) Not at all — (4) Moderately — (7) Completely			

Results. After excluding participants who failed the AI and attention checks, the final sample consisted of 186 participants (43% male, 55.4% female, 1.6% preferred not to say). We conducted contrast one-way ANOVAs with planned contrasts to examine differences across hierarchy levels in the participant ratings of noun representativeness, verb-noun pair adequacy, and verb accuracy. For the noun representativeness, level 6 (Mean = 5.39) yielded the highest

mean rating, followed by level 4 (Mean = 4.61), level 5 (Mean = 4.59), and level 3 (Mean = 4.22). Level 6 differed significantly from all other levels ($p = .000$), whereas level 3, level 4, and level 5 did not significantly differ from one another ($p = .316$ and $p = .380$, respectively). For the verb-noun pair adequacy, level 6 again yielded the highest mean value (Mean = 5.19) and differed significantly from all other levels (level 3: Mean = 3.95, $p = .000$; level 4: Mean = 4.19, $p = .000$; level 5: Mean = 4.19, $p = .000$). No significant difference between level 3, 4 and 5 ($ps = 1.000$). For the accuracy of verb, mean ratings were consistently high across all levels ($M = 6.02, 6.14, 6.16$, and 6.07 for levels 3–6, respectively), with no significant differences emerged ($ps = 1.000$). Taken together, these results support the selection of level 6 for the main analysis, as it yields significantly higher representativeness rating for both nouns and verb-noun pairs relative to level 3 through 5, while verb accuracy remains uniformly high across all levels.

Second, we recruited 41 participants from Prolific to validate the verb-noun pair matching. The final sample includes 30 participants (average age = 32.86) after excluding participants who failed AI check and technology description fluency questions. Each participant was randomly assigned to read 10 matched or non-matched patent and task descriptions, then answer the following question. The correct classification rate is 72.59% among 300 task-patent descriptions.

“Many AI patents describe general-purpose technologies that were not originally designed for any specific marketing task but could plausibly be adapted or repurposed to support one. With this in mind: does the AI technology described in the patent have plausible *potential applicability* to the task described in the job posting. That is, could it reasonably serve as an underlying capability or technical building block for performing or assisting with this task, even if substantial adaptation, integration, or recombination with other technologies would be required?”

Part 2. Depth and Breadth Validation

To validate the distinction between depth and breadth, we conducted both a qualitative study and one quantitative study to establish validity through practitioner input and to test whether the two constructs are empirically different. For the qualitative study, we invited 8 practitioners with at least 7 years of professional experience to respond to five open-ended questions designed across whether the conceptual distinction between depth and breadth aligns with practitioners' understanding of AI's impact on marketing roles (see Table E2 for the question list and Table E3 for survey sample description).

Table E2. Depth and Breadth Questions for Practitioners

1. In your view, is there a difference between a job having many AI tools relevant to it and AI being able to perform a large share of that job's tasks? Please explain.
2. Can a marketing-related job have many AI tools relevant to it, but still remain relatively secure from replacement? Why or why not?
3. Can AI make someone in a marketing-related role more productive or effective without making the role more replaceable? Why or why not?
4. What kinds of marketing-related jobs seem most vulnerable to AI replacement, and what makes them vulnerable?
5. What kinds of marketing jobs seem most likely to benefit from AI through higher productivity or effectiveness, even if the job is not replaced?

Table E3. Survey Sample Description

#	Job Description	Tenure (Years)	Industry
1	Marketing VP	More than 8 years	Food
2	Founder	4 – 7 years	AI
3	Lifecycle Marketing Manager	More than 8 years	E-commerce
4	CEO	More than 8 years	Ecological environment
5	Marketing Manager	4 – 7 years	Ecological environment
6	Chairman	More than 8 years	Trade and Logistics
7	Product Manager	More than 8 years	Technology
8	Key Account Manager	4 – 7 years	Marketing

For the quantitative study, we conducted a 2x2 between-subject design and recruited 200 participants whose primary language is English from Prolific, an online crowdsourcing platform. Each participant was randomly assigned to one of four conditions defined by depth level (low vs. high) and breadth level (low vs. high), read a marketing job description, and completed measures assessing perceived AI replacement likelihood, productivity enhancement, augmentation, and reskilling needs, along with manipulation checks for depth and breadth. Table E4 presents the details of the study design and measures.

Table E4. Depth and Breadth (Quantitative Study)

Condition A (low breadth, low depth)	There are only <i>a few AI technologies</i> available for this marketing analyst's job. These technologies apply to only a limited set of the analyst's tasks, and each can assist with only <i>a small portion</i> of the overall work. Most of the tasks in the job still require substantial human judgment and effort.
Condition B (high breadth, low depth)	There are <i>many different AI technologies</i> available for this marketing analyst's job. These technologies apply to only a limited set of the analyst's tasks, and each can assist with only <i>a small portion</i> of the overall work. Most of the tasks in the job still require substantial human judgment and effort.
Condition C (low breadth, high depth)	There are only <i>a few AI technologies</i> available for this marketing analyst's job. However, the AI technologies that do apply can perform <i>a large share</i> of the overall work. As a result, a substantial portion of the job can be completed by limited AI technologies rather than by a human employee.
Condition D (high breadth, high depth)	There are <i>many different AI technologies</i> available for this marketing analyst's job. These technologies apply across several parts of the analyst's work, and together they can perform <i>a large share</i> of the overall work. As a result, a substantial portion of the job can be completed by many AI technologies rather than by a human employee.
Manipulation	Many different AI technologies are relevant to this job. This job exposes workers to many different AI technologies. AI can perform a large share of the tasks in the job. AI has deeply penetrated the task structure of this job. (1) Strongly disagree ----- (4) Neutral ---- (7) Strongly agree

Results. After excluding participants who failed the AI and attention checks, the final sample consisted of 164 participants (39.6% male, 59.1% female, 1.3% preferred not to say, $M_{age} = 38.90$). Manipulation checks confirmed that both manipulations were effective. For breadth, participants in the high condition perceived significantly greater breadth than those in the low condition on both items (“Many different AI technologies are relevant to this job”: $M_{high} = 5.89$, $SD = 1.136$ vs. $M_{low} = 3.23$, $SD = 2.050$; $t(162) = -10.214$, $p = .000$; “This job exposes workers to many different AI technologies”: $M_{high} = 5.61$, $SD = 1.317$ vs. $M_{low} = 3.21$, $SD = 1.914$; $t(162) = -9.303$, $p = .000$). For depth, participants in the high condition likewise perceived significantly greater depth than those in the low condition on both items (“AI can perform a large share of the tasks in the job”: $M_{high} = 6.10$, $SD = 1.038$ vs. $M_{low} = 3.34$, $SD = 1.913$; $t(162) = -11.468$, $p = .000$; “AI has deeply penetrated the task structure of this job”: $M_{high} = 5.46$, $SD = 1.335$ vs. $M_{low} = 3.50$, $SD = 1.752$; $t(162) = -11.468$, $p = .000$).

Web Appendix F: Descriptive Statistics by Occupation Groups and Industries

Figures F1 and F2 show that AI exposure also varies across occupation groups and industries. Communication and public relations exhibit low values on both dimensions, consistent with the importance of relational judgment, narrative framing, and contextual interpretation. Customer-service-related roles exhibit high values on both dimensions, consistent with the codifiable and repeatable structure of many service interactions. Industry patterns also reveal meaningful heterogeneity. Retail trade, transportation and warehousing, public administration, accommodation and food services, and finance and insurance are concentrated in the high-depth, high-breadth quadrant, while entertainment and recreation, educational services, construction, and information display lower exposure on both dimensions.

Figure F1. AI Depth and Breadth by Occupation Group





Figure F2. AI Breadth and Depth by Industries



Web Appendix G: Skill Classification

In the job posting dataset, Lightcast provides a list of skills collected from job postings and classifies each skill into a unique taxonomy of 31 categories with 3-tier hierarchy, using a combination of machine-learning and in-house experts validation (full list at <https://lightcast.io/open-skills/categories>). This taxonomy is provided directly by Lightcast, and we do not modify the underlying skill-to-category assignment.

From these 31 categories, we selected those most relevant to marketing work, and they are marketing analytics, AI skills, customer-management skills, sales skills, brand management skills, marketing operations skills, and integrated marketing communication skills. Specifically, AI skills correspond directly to Lightcast's AI/Machine Learning category. Sales skills and brand management skills correspond to Lightcast's Sales and Brand Management categories, respectively. Customer-management skills correspond to Lightcast's Customer and Client Support category. Marketing operations skills correspond to Lightcast's Marketing Strategy and Techniques category. We aggregate two categories, which are Lightcast's Market Analysis and Marketing Software categories, into marketing analytics skills. IMC skills correspond to seven Lightcast categories, which are Advertising, Online Advertising, Digital Marketing, Promotions and Campaigns, Public Relations, Social Media, and Web Analytics and SEO categories, into IMC skills.

Web Appendix H: Control Function First-Stage Results

Table H1. Job Hiring - First Stage

Variables	Model 1 (Applicability Breadth)			Model 2 (Substitution Depth)		
	Estimate	S.E.	p-value	Estimate	S.E.	p-value
Industry's peer patent	-0.005	0.003	0.054	0.000	0.002	0.842
Datacenter count	-0.010	0.007	0.158	0.014	0.006	0.011
Mkt analytics skill	0.203	0.019	0.000	0.155	0.017	0.000
AI skill	0.000	0.033	0.995	-0.115	0.025	0.000
Customer-relationship skill	0.288	0.010	0.000	0.107	0.006	0.000
Sales skill	-0.156	0.005	0.000	-0.009	0.004	0.039
Brand marketing skill	-0.396	0.022	0.000	-0.118	0.017	0.000
Marketing operations skill	-0.244	0.010	0.000	-0.090	0.008	0.000
IMC skill	-0.075	0.010	0.000	-0.029	0.007	0.000
Num responsibilities	-0.013	0.001	0.000	-0.002	0.001	0.000
Min years of experience	-0.048	0.001	0.000	-0.032	0.001	0.000
Education	0.024	0.002	0.000	0.004	0.002	0.022
Disclosure Law	0.002	0.003	0.384	-0.002	0.002	0.363
Occupation FE	Yes			Yes		
MSA FE	Yes			Yes		
Time FE	Yes			Yes		
Firm FE	Yes			Yes		
Observations	1,549,602			1,549,602		
R ²	.429			.724		
F-statistics for IV	10.39 p= .000			9.65 p= .000		

Table H2. Salary - First Stage

Variables	Model 1 (Applicability Breadth)			Model 2 (Substitution Depth)		
	Estimate	S.E.	p-value	Estimate	S.E.	p-value
Industry's peer patent	-0.002	0.0007	0.003	0.005	0.001	0.000
Datacenter count	-0.023	0.006	0.000	0.011	0.006	0.059
Mkt analytics skill	0.197	0.014	0.000	0.329	0.015	0.000
AI skill	0.081	0.047	0.081	-0.064	0.020	0.002
Customer-relationship skill	0.243	0.008	0.000	0.107	0.005	0.000
Sales skill	-0.124	0.004	0.000	0.052	0.004	0.000
Brand Management skill	-0.426	0.017	0.000	-0.168	0.016	0.000
Marketing operations skill	-0.234	0.009	0.000	-0.122	0.008	0.000
IMC skill	-0.052	0.009	0.000	0.022	0.007	0.002
Num responsibilities	-0.022	0.001	0.000	-0.009	0.000	0.000
				5		
Min years of experience	-0.048	0.001	0.000	-0.039	0.001	0.000
Education	0.029	0.002	0.000	-0.004	0.001	0.011
Disclosure Law	-0.002	0.002	0.248	-0.005	0.002	0.016
Occupation FE		Yes			Yes	
MSA FE		Yes			Yes	
Time FE		Yes			Yes	
Firm FE		Yes			Yes	
Observations		2,297,723			2,297,723	
R ²		.383			.624	
F-statistics for IV		3.27 p=.038			2.86 p=.057	

Web Appendix I: Full Results

Table I1. Full Results of Job Hiring Model

	Model 1			Model 2			Model 3			Model 4			Model 5		
	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value
Applicability Breadth	1.307	.238	.000	1.438	.240	.000	1.445	.240	.000	1.442	.240	.000	1.448	.240	.000
Substitution Depth	-2.214	.430	.000	-2.326	.431	.000	-2.330	.431	.000	-2.323	.431	.000	-2.329	.431	.000
Breadth × Depth	—	—	—	-.060	.002	.000	-.061	.002	.000	-.064	.002	.000	-.064	.002	.000
Marketing Analytics × Breadth	—	—	—	—	—	—	-.079	.027	.004	—	—	—	-.091	.028	.001
Customer-management × Breadth	—	—	—	—	—	—	.003	.007	.619	—	—	—	.007	.007	.285
AI × Breadth	—	—	—	—	—	—	-.178	.056	.001	—	—	—	-.134	.061	.027
Sales × Breadth	—	—	—	—	—	—	-.003	.004	.474	—	—	—	.012	.005	.011
Brand Management × Breadth	—	—	—	—	—	—	.080	.032	.013	—	—	—	.098	.034	.004
Marketing Operations × Breadth	—	—	—	—	—	—	.111	.013	.000	—	—	—	.100	.013	.000
IMC × Breadth	—	—	—	—	—	—	-.082	.010	.000	—	—	—	-.067	.011	.000
Marketing Analytics × Depth	—	—	—	—	—	—	—	—	—	.008	.021	.719	.029	.022	.191
Customer-management × Depth	—	—	—	—	—	—	—	—	—	-.009	.010	.389	-.012	.011	.273
AI × Depth	—	—	—	—	—	—	—	—	—	-.077	.030	.009	-.051	.032	.114
Sales × Depth	—	—	—	—	—	—	—	—	—	-.053	.004	.000	-.056	.004	.000

Brand Management × Depth	—	—	—	—	—	—	—	—	—	-.033	.020	.096	-.054	.021	.011
Marketing Operations × Depth	—	—	—	—	—	—	—	—	—	.053	.010	.000	.027	.010	.006
IMC × Depth	—	—	—	—	—	—	—	—	—	-.059	.007	.000	-.043	.007	.000
Marketing Analytic	.131	.045	.004	.122	.045	.007	.110	.046	.016	.122	.045	.007	.109	.046	.017
Customer-management	.146	.043	.001	.131	.044	.003	.129	.044	.003	.132	.044	.003	.129	.044	.003
AI Skill	-.087	.070	.216	-.095	.070	.178	-.159	.072	.027	-.166	.072	.020	-.191	.073	.008
Sales Skill	.340	.035	.000	.353	.035	.000	.353	.035	.000	.351	.035	.000	.353	.035	.000
Brand Management	.206	.067	.002	.237	.067	.000	.254	.068	.000	.233	.068	.001	.252	.068	.000
Marketing Operations	.216	.037	.000	.236	.038	.000	.253	.038	.000	.243	.038	.000	.256	.038	.000
IMC	-.011	.015	.469	-.009	.015	.547	-.025	.015	.110	-.027	.016	.087	-.035	.016	.024
Number of Responsibilities	-.008	.002	.001	-.008	.002	.001	-.008	.002	.001	-.008	.002	.001	-.008	.002	.001
Min. Years of Experience	-.032	.008	.000	-.030	.008	.000	-.029	.008	.000	-.029	.008	.000	-.029	.008	.000
Education	-.125	.005	.000	-.127	.005	.000	-.127	.005	.000	-.127	.005	.000	-.128	.005	.000
Disclosure Law	.039	.004	.000	.038	.004	.000	.038	.004	.000	.038	.004	.000	.038	.004	.000
Control Function	-1.296	.238	.000	-1.400	.240	.000	-1.406	.240	.000	-1.402	.240	.000	-1.411	.240	.000
Residual 1	2.225	.430	.000	2.308	.431	.000	2.312	.431	.000	2.319	.431	.000	2.327	.431	.000
Control Function															
Residual 2															
Occupation, MSA, Time, Firm Fixed Effects	YES			YES			YES			YES			YES		
Observations	1,493,148			1,493,148			1,493,148			1,493,148			1,493,148		
R ²	.271			.271			.271			.271			.271		

Table I2. Full Results of Salary Model

	Model 1			Model 2			Model 3			Model 4			Model 5		
	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value
Applicability Breadth	-2.429	.358	.000	-2.429	.373	.000	-2.438	.373	.000	-2.452	.373	.000	-2.446	.373	.000
Substitution Depth	-.887	.385	.021	-.887	.377	.019	-.890	.377	.018	-.913	.377	.015	-.902	.377	.017
Breadth × Depth	—	—	—	-.007	.001	.000	-.007	.001	.000	-.006	.001	.000	-.007	.001	.000
Marketing Analytics × Breadth	—	—	—	—	—	—	-.072	.018	.000	—	—	—	-.064	.019	.001
Customer-management × Breadth	—	—	—	—	—	—	.089	.002	.000	—	—	—	.077	.002	.000
AI × Breadth	—	—	—	—	—	—	.207	.018	.000	—	—	—	.252	.014	.000
Sales × Breadth	—	—	—	—	—	—	.012	.002	.000	—	—	—	.007	.002	.004
Brand Management × Breadth	—	—	—	—	—	—	.020	.014	.146	—	—	—	.050	.014	.000
Marketing Operations×Breadth	—	—	—	—	—	—	.027	.005	.000	—	—	—	.040	.005	.000
IMC × Breadth	—	—	—	—	—	—	.117	.004	.000	—	—	—	.099	.004	.000
Marketing Analytics × Depth	—	—	—	—	—	—	—	—	—	-.043	.016	.005	-.022	.016	.179
Customer-management × Depth	—	—	—	—	—	—	—	—	—	.110	.004	.000	.079	.004	.000
AI × Depth	—	—	—	—	—	—	—	—	—	-.116	.020	.000	-.163	.019	.000
Sales × Depth	—	—	—	—	—	—	—	—	—	.026	.002	.000	.022	.002	.000
IMC × Depth	—	—	—	—	—	—	—	—	—	.068	.005	.000	.055	.005	.000
Brand Management × Depth	—	—	—	—	—	—	—	—	—	.026	.014	.063	-.075	.014	.000
Marketing Operations × Depth	—	—	—	—	—	—	—	—	—	.019	.005	.000	-.044	.006	.000
Marketing Analytics	1.499	.183	.000	1.498	.188	.000	1.486	.187	.000	1.506	.187	.000	1.491	.187	.000
Customer-management	.338	.120	.005	.338	.125	.007	.318	.125	.011	.318	.125	.011	.303	.125	.015
AI Skill	1.768	.033	.000	1.768	.040	.000	1.836	.040	.000	1.643	.047	.000	1.675	.046	.000
Sales Skill	-.218	.033	.000	-.218	.034	.000	-.217	.034	.000	-.220	.034	.000	-.218	.034	.000
Brand Management	-1.302	.204	.000	-1.302	.212	.000	-1.295	.212	.000	-1.307	.212	.000	-1.304	.212	.000

Marketing Operations	-.624	.121	.000	-.624	.126	.000	-.621	.126	.000	-.630	.126	.000	-.629	.126	.000
IMC	-.437	.016	.000	-.437	.016	.000	-.418	.016	.000	-.417	.016	.000	-.405	.016	.000
Number of Responsibilities	-.058	.011	.000	-.058	.011	.000	-.058	.011	.000	-.059	.011	.000	-.058	.011	.000
Min. Years of Experience	.253	.030	.000	.253	.031	.000	.252	.031	.000	.251	.031	.000	.252	.031	.000
Education	.016	.010	.099	.016	.010	.111	.016	.010	.104	.017	.010	.089	.017	.010	.091
Disclosure Law	-.002	.003	.432	-.002	.003	.437	-.003	.003	.426	-.002	.003	.429	-.002	.003	.436
Control Function Residual Breadth	2.395	.357	.000	2.395	.373	.000	2.391	.373	.000	2.417	.373	.000	2.400	.373	.000
Control Function Residual Depth	.875	.385	.023	.875	.377	.020	.878	.377	.020	.888	.377	.019	.881	.377	.019
Occupation, MSA, Time, Firm Fixed Effects		YES		YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations		2,297,723			2,297,723			2,297,723			2,297,723			2,297,723	
R ²		.698			.697			.698			.698			.698	

Web Appendix J: Robustness Check

Table J1. Robustness Checks: Job Hiring

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	No Controls	Importance Score	Interacted FE	Interacted FE +controls	Eloundou measure	Full Sample	Full Sample + controls	Monthly	Log posting
AI Breadth	0.018(0.000)	0.020(0.000)	0.025(0.000)	0.050(0.000)	0.036(0.000)	0.018(0.000)	0.024(0.000)	0.032(0.000)	0.019(0.000)
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.0008)
AI Depth	-0.006(0.000)	-0.010(0.000)	-0.009(0.000)	-0.025(0.000)	-0.013(0.000)	-0.001(0.016)	-0.004(0.016)	-0.010(0.000)	-0.006(0.000)
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Mkt Analytics Skill Ratio		0.052(0.018)		0.105(0.001)	0.115(0.000)		0.040(0.000)	0.118(0.000)	0.027(0.025)
		(0.022)		(0.032)	(0.022)		(0.010)	(0.022)	(0.012)
AI Skill Ratio		0.169(0.001)		0.101(.095)	0.198(0.000)		0.121(0.000)	0.270(0.000)	0.081(0.002)
		(0.049)		(0.061)	(0.045)		(0.020)	(0.045)	(0.026)
Customer Skill Ratio		0.286(0.001)		0.352(0.000)	0.262(0.000)		0.226(0.000)	0.215(0.000)	0.152(0.000)
		(0.010)		(0.015)	(0.010)		(0.006)	(0.009)	(0.012)
Sales Skill Ratio		0.157(0.000)		0.224(0.000)	0.141(0.000)		0.114(0.000)	0.117(0.000)	0.079(0.000)
		(0.006)		(0.008)	(0.006)		(0.003)	(0.005)	(0.003)
Brand Management Skill Ratio		-0.044(0.088)		- 0.051(0.162)	0.031(0.127)		-0.023(0.014)	0.072(0.000)	-0.031(0.025)
		(0.026)		(0.037)	(0.020)		(0.009)	(0.019)	(0.014)
Marketing Operations Skill Ratio		0.102(0.000)		0.137(0.000)	0.162(0.000)		0.070(0.000)	0.107(0.000)	0.043(0.000)
		(0.012)		(0.017)	(0.012)		(0.005)	(0.010)	(0.006)
IMC Skill Ratio		-0.048(0.000)		-0.102(0.000)	-0.063(0.000)		0.037(0.000)	-0.060(0.000)	-0.028(0.000)
		(0.011)		(0.016)	(0.008)		(0.007)	(0.008)	(0.006)
Num. Responsibilities		-0.020(0.000)		-0.027(0.000)	-0.018(0.000)		-0.010(0.000)	-0.014(0.000)	-0.010(0.000)
		(0.001)		(0.001)	(0.001)		(0.0001)	(0.001)	(0.0005)

Min Years Experience	-0.023(0.000)			-0.026(0.000)	-0.015(0.000)		-0.013(0.000)	-0.009(0.000)	-0.014(0.000)
	(0.001)			(0.001)	(0.001)		(0.0001)	(0.001)	(0.0005)
Education	-0.103(0.000)			-0.125(0.000)	-0.094(0.000)		-0.079(0.000)	-0.080(0.000)	-0.048(0.000)
	(0.003)			(0.004)	(0.002)		(0.001)	(0.002)	(0.001)
Disclosure Law	0.047(0.004)			0.035(0.000)	0.028(0.000)		0.125(0.000)	0.010(0.000)	0.030(0.000)
	(0.004)			(0.006)	(0.006)		(0.003)	(0.005)	(0.002)
Exposure × Penetration	-0.030(0.000)	-0.031(0.000)	-0.044(0.000)	-0.082(0.000)	-0.053(0.000)	-0.024(0.000)	-0.035(0.000)	-0.049(0.000)	-0.028(0.000)
	(0.001)	(0.001)	(0.002)	(0.003)	(0.003)	(0.001)	(0.001)	(0.002)	(0.001)
Eloundou					-0.492(0.000)				
					(0.041)				
Occ, MSA, Time, Firm Fixed Effects	YES	YES	Occ×MSA, Time ×Firm	Occ×MSA, Time ×Firm	YES	YES	YES	YES	YES
Observations	2,841,768	1,519,911	2,305,477	1,196,869	1,457,054	10,823,419	5,672,109	1,710,596	1,493,148
R ²	0.230	0.268	0.326	0.372	0.153	0.197	0.1004	0.144	0.293

Notes: Dependent variable is standardized total job postings (1-quarter lead), coefficients are reported with standard errors in parentheses below each estimate and p-values in parentheses to the right. Eloundou's measure is significant in the model without our measures. The changes in the number of observations are due to missing values.

Table J2. OLS Robustness Checks: Salary

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	No Controls	Importance Score	Interacted FE	Interacted FE + Controls	Eloundou measure	Disclosure only	Non-disclosure only	Log Salary
AI Breadth	-0.028(0.000)	-0.034(0.000)	-0.032(0.000)	-0.034(0.000)	-0.019(0.000)	-0.036(0.000)	-0.033(0.000)	-0.020(0.000)
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.0003)
AI Depth	-0.019(0.000)	-0.024(0.000)	-0.021(0.000)	-0.012(0.000)	-0.098(0.000)	-0.017(0.000)	-0.011(0.000)	-0.001(0.004)
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.0004)
		0.519(0.000)		0.793(0.000)	0.735(0.000)	0.682(0.000)	0.739(0.000)	0.382(0.000)

Mkt Analytics Skill Ratio		(0.016)		(0.017)	(0.017)	(0.032)	(0.020)	(0.007)
Customer Skill Ratio		-0.284(0.000)		-0.338(0.000)	-0.339(0.000)	-0.377(0.000)	-0.330(0.000)	-0.253(0.000)
		(0.003)		(0.004)	(0.004)	(0.008)	(0.004)	(0.002)
AI Skill Ratio		1.512(0.000)		1.630(0.000)	1.625(0.000)	1.657(0.000)	1.498(0.000)	0.471(0.000)
		(0.034)		(0.035)	(0.035)	(0.055)	(0.044)	(0.012)
Sales Skill Ratio		0.026(0.000)		0.034(0.000)	0.029(0.000)	-0.87(0.000)	0.067(0.000)	0.014(0.000)
		(0.003)		(0.003)	(0.003)	(0.006)	(0.004)	(0.001)
Brand Management Skill Ratio		-0.150(0.000)		-0.134(0.000)	-0.131(0.000)	-0.137(0.000)	-0.139(0.000)	-0.056(0.000)
		(0.014)		(0.014)	(0.014)	(0.025)	(0.017)	(0.001)
Marketing Operations Skill Ratio		0.008(0.000)		0.042(0.000)	0.038(0.000)	0.028(0.049)	0.041(0.000)	0.032(0.000)
		(0.007)		(0.007)	(0.007)	(0.014)	(0.008)	(0.003)
IMC Skill Ratio		-0.348(0.000)		-0.331(0.000)	-0.331(0.000)	-0.343(0.000)	-0.326(0.000)	-0.132(0.000)
		(0.007)		(0.007)	(0.007)	(0.014)	(0.004)	(0.003)
Num. Responsibilities		0.003(0.000)		-0.002(0.000)	0.002(0.000)	0.004(0.000)	0.003(0.000)	0.006(0.000)
		(0.001)		(0.001)	(0.001)	(0.002)	(0.001)	(0.0003)
Min. Years of Experience		0.366(0.000)		0.401(0.000)	0.401(0.000)	0.467(0.000)	0.376(0.000)	0.168(0.000)
		(0.001)		(0.001)	(0.001)	(0.002)	(0.001)	(0.0003)
Education		-0.045(0.000)		-0.051(0.000)	-0.051(0.000)	-0.022(0.000)	-0.058(0.000)	-0.015(0.072)
		(0.001)		(0.002)	(0.001)	(0.003)	(0.002)	(0.001)
Disclosure Law		0.08(0.000)		0.006(0.004)	0.006(0.005)	-	-	0.001(0.000)
		(0.002)		(0.002)	(0.002)			(0.001)
Exposure × Penetration	-0.007(0.000)	0.009(0.000)	-0.005(0.000)	-0.007(0.000)	-0.002(0.031)	-0.008(0.002)	-0.005(0.000)	-0.006(0.000)
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Eloundou					-0.098(0.000)			
					(0.002)			

Occ, MSA, Time, Firm Fixed Effects	YES	YES	Occ×MSA, Date×Firm	YES	YES	YES	YES	YES
Observations	4,705,945	2,297,723	3,929,521	2,297,723	2,290,664	540,665	1,739,140	2,297,723
R ²	0.628	0.718	0.714	0.697	0.698	0.757	0.686	0.771

Notes: Dependent variable is standardized salary, coefficients are reported with standard errors in parentheses below each estimate and p-values in parentheses to the right. Eloundou's measure is significant in the model without our measures. The changes in the number of observations are due to missing values.

Reference

- Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller (2025), “Artificial Intelligence and the Labor Market,” Working Paper Series, Working Paper, National Bureau of Economic Research.
- Martin, Andre and Tarun Kushwaha (2024), “Can Words Speak Louder Than Actions? Using Top Management Teams’ Language to Predict Myopic Marketing Spending,” *Journal of Marketing*, 00222429241244804.
- Miller, George A. (1995), “WordNet: a lexical database for English,” *Commun. ACM*, 38 (11), 39–41.
- Montani, Ines, Matthew Honnibal, Matthew Honnibal, Adriane Boyd, Sofie Van Landeghem, and Henning Peters (2023), “explosion/spaCy: v3.7.2: Fixes for APIs and requirements,” Zenodo.
- Webb, Michael (2019), “The Impact of Artificial Intelligence on the Labor Market,” SSRN Scholarly Paper, Rochester, NY.