Workforce Planning in the Age of AI

Example Co.

Pierre Bouquet
Nicolo' Piergiovanni Bagnoli
Yossi Sheffi

November 2025

MIT Center for
Transportation & Logistics

Introduction

In 2024, ExampleState experienced the fastest population growth in the nation, placing unprecedented demands on the Company (ExampleCo) to deliver and maintain nearly 80,000 center-line miles of roadway (Company, 2024). Continued population growth, supply chain disruptions, and rising costs highlight the need for ExampleCo to operate more efficiently. To address these challenges, ExampleCo must analyze its internal operations in detail and strategically adopt emerging technologies to maintain economic resilience, support national commerce, and meet increasing mobility needs (Bureau of Transportation Statistics, 2023; Sheffi, 2025).

Artificial intelligence (AI) has quickly become a general-purpose technology with the potential to transform administrative, engineering, and operational workflows (Eloundou et al., 2024). Transportation and infrastructure agencies already use AI to automate supplier negotiations (Hoek et al., 2022), optimize supply chain management (Rolf et al., 2022), and improve customer service (Kang & Choi, 2023). As these applications advance, AI presents ExampleCo with a significant opportunity to reduce administrative workload, streamline processes, and allow staff to focus on the agency's core mission of planning, building, and maintaining the transportation network for nearly 30 million Texans.

This report addresses ExampleCo leadership's key questions about AI's impact on the workforce.

Where should ExampleCo deploy AI, what benefits are expected, and how will it affect the workforce?

To answer these questions, we introduce a new methodology that breaks jobs into their basic units: tasks. For each task, we evaluate its cost and assess to what extent it can be accelerated by an AI-based system. This task-based framework allows for a precise assessment of AI's impact across roles, departments, and the organization.

This research was conducted at the MIT Center for Transportation & Logistics, with a focus on artificial intelligence, labor economics, and supply chain management. The report provides data-driven insights to support ExampleCo's strategic decisions regarding AI's impact and investment. Specifically, the report will:

A. Map ExampleCo workforce to a task-level representation, allocating wage bill and FTE equivalents to individual activities using O*NET and an AI-powered crosswalk framework.

B. Quantify the exposure of each task and job to AI, using a dynamic news-based exposure metric that captures how emerging technologies affect specific activities over time.

C. Simulate targeted deployment scenarios, estimating how productivity-enhancing AI tools – such as Microsoft Copilot – could reduce wage bill costs and support reinvestment in skills, operations, and employee wellbeing.

A.1 ExampleCo Workforce Review
People, wage bill and growth

13,081
Number of workers in ExampleCo.

ExampleCo Employment

- In 2025, ExampleCo employed around 13 thousand persons.
- This number represents a 13.19% total growth from January 2018.

$1B
Wage bill of ExampleCo.

ExampleCo Wage Bill

- In 2025, ExampleCo’s wage bill (sum of all employee salaries) was just over one billion dollars
- The wage bill increased by 56.65% since 2018 – when it was $639M.

-0.91%
2024 to 2025 headcount change.

ExampleCo Employment

- Employee count declined in 2025 by 0.91%, marking the first contraction since 2022, when headcount fell by 0.10%.

Exhibit 1 – ExampleCo Wage bill, employee count, and respective YoY rates.

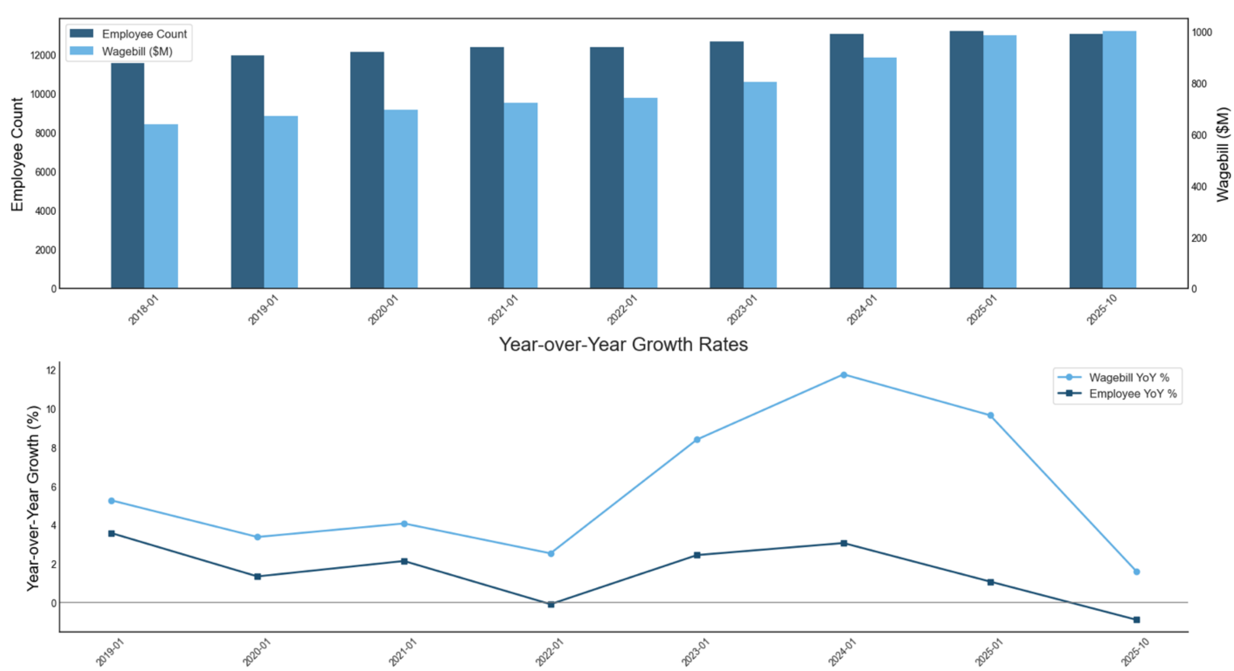

Source: ExampleCo 2025

A.2 | Tenure
Tenure and trends in the workforce

60
Total districts.

Districts in ExampleCo

- The wage bill of the largest district is $82.8M, with more than 1200 FTEs.

1095
Total departments.

Departments/Divisions in ExampleCo

- The wage bill of the largest is just under $9M, with over 140 FTEs.

Workforce changes within departments and across ExampleCo geographies reveal significant internal shifts in department distribution.

From 2018 to 2025, ExampleCo's organizational structure remains stable at the district level, while departments experience notable changes. District and division counts fluctuate only between 60 and 61, with two districts opening and two closing, resulting in no net structural change.

In contrast, departments show significant activity, increasing from 1,022 to 1,095, a net gain of 73. This change results from 136 department openings and 80 closures, reflecting ongoing reconfiguration and expansion to meet operational needs. The data suggest that while ExampleCo's geographic structure is stable, its departmental organization remains dynamic to support evolving priorities and workforce requirements.

Exhibit 2 – ExampleCo largest 15 Departments and Districts, the box includes the growth from 2018 to 2025 and the $ value represents the Wage Bill.

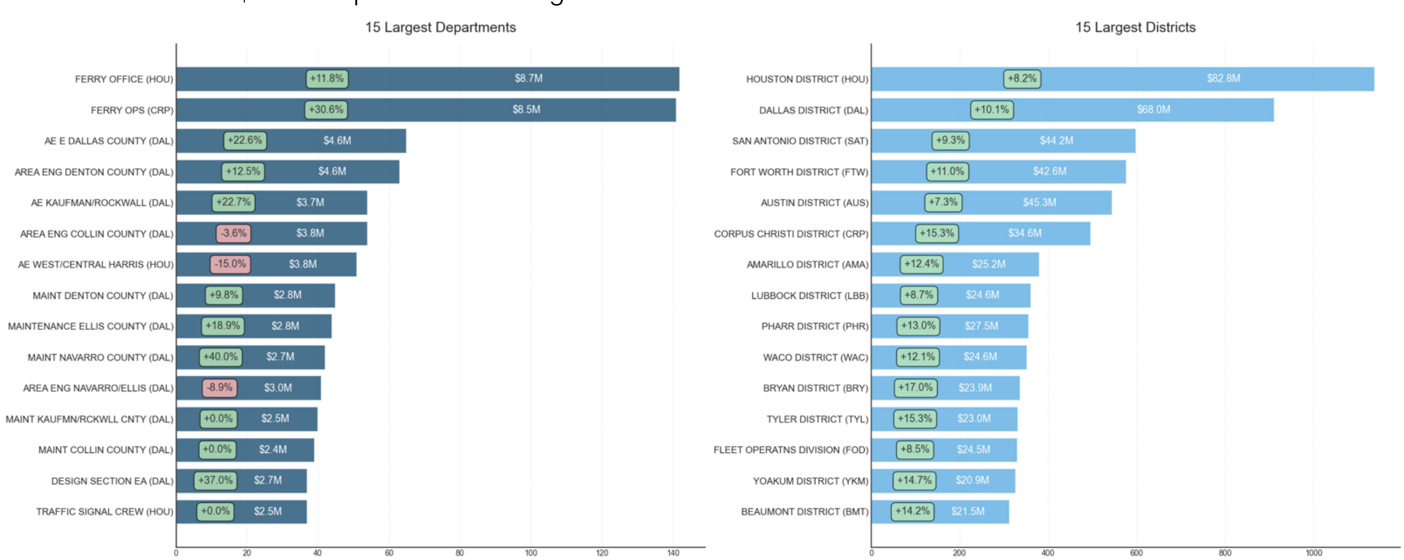

Source: ExampleCo 2025

A.3 Tenure

Tenure and trends in the workforce

-30.8%

Decrease of "fewer than 2 years" tenure employees at ExampleCo since 2024.

Fall in "fewer than 2 years" employed category since 2024

- Since the start of 2024, the "fewer than 2 years" of tenure class of employees has drastically been reduced. This, along with the steady decrease of seniors, at the ExampleCo is a trend to be investigated.

Analysis of ExampleCo employees by five tenure classes reveals clear and notable trends.

ExampleCo's tenure distribution shows notable changes from 2018 to 2025. Tenure bands follow SAO workforce reporting conventions and are based on completed years of service. The agency has experienced a steady decline in employees with 15 or more years of service, indicating ongoing retirements and departures among its most experienced staff. In contrast, mid-career groups, especially those with 5 to 9 years of service, have grown consistently and now form one of the most stable workforce segments. The 10 to 14-year band has also rebounded after an initial decline, reflecting a maturing group of employees who joined in the late 2010s.

The "Fewer than 2 years" cohort has shown significant volatility. Following rapid growth in the early 2020s, likely due to increased post-pandemic hiring, this group has recently declined sharply. This trend may reflect reduced hiring, higher turnover among new staff, or slower onboarding. The decrease is important, as early-tenure employees often reflect organizational health. A reduced inflow of new talent, combined with fewer long-tenured employees, may signal a tightening workforce pipeline that ExampleCo should monitor to maintain long-term capacity and institutional knowledge.

Exhibit 3 – Trends in ExampleCo tenure (2018-2025).

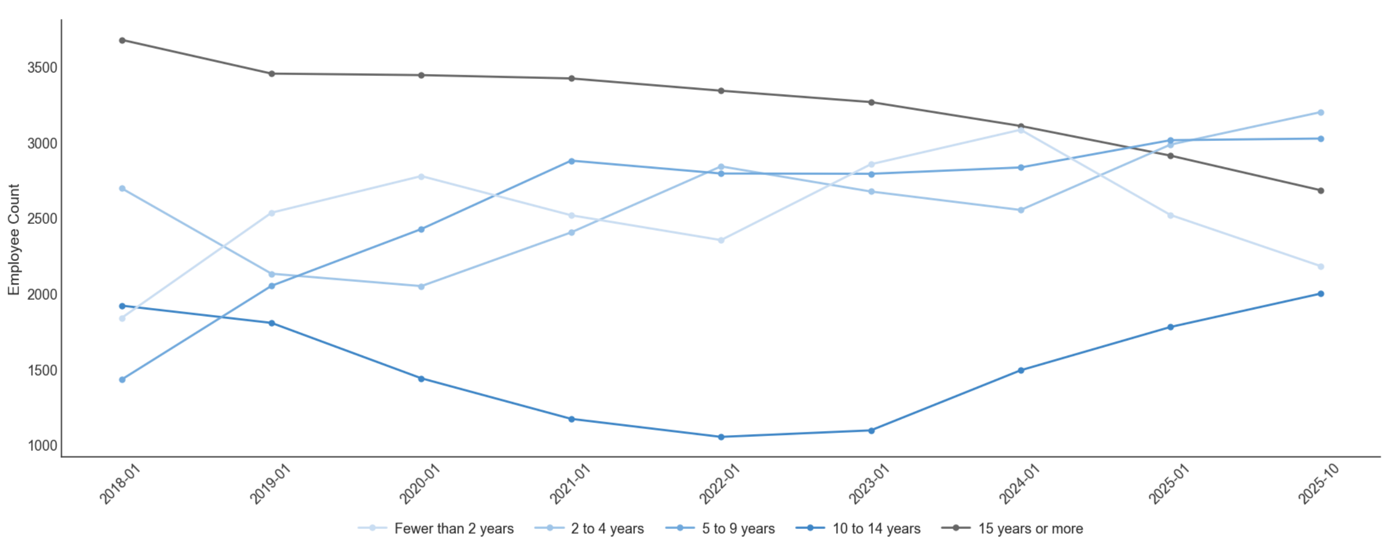

Source: ExampleCo 2025

B.1 Methodology
Jobs as Tasks

A common framework in labor-economics research represents jobs as collections of tasks (Autor et al., 2003).

To implement this framework, we utilize the U.S. Department of Labor's Occupational Information Network (National Center for O*NET Development, 2025), which reports a detailed list of tasks based on the frequency with which workers perform them for each job and categorizes them by sector.

We treat task frequency as a proxy for the allocation of a worker's time dedicated to each task. This approach assumes that each task requires a comparable amount of time for a 'typical' worker in that occupation (Martin, 2022; Brandes & Wattenhofer, 2016).

Full-time equivalent (FTE)

An FTE converts task time shares into the equivalent number of full-time workers. For example, if a task occupies 25% of a worker's time, that equals 0.25 FTE. Summing across all employees in an occupation tells us how many full-time workers would be required to cover that task.

Task Annual Wage Bill

To evaluate the annual wage bill of each task, we multiply its FTE total by the occupation's average yearly salary. For instance, if a task represents 10 FTEs and the average salary is $80,000, the task's annual wage bill is $800,000.

Task Wage Bill and FTE Across Occupations

Many tasks appear in multiple occupations. By aggregating a task's FTEs and wage bills across all relevant jobs, we can estimate the total labor cost dedicated to that task workforce-wide.

Exhibit 4 – Example with the "Graders and Sorters, Agricultural Products" job represented as a set of tasks.

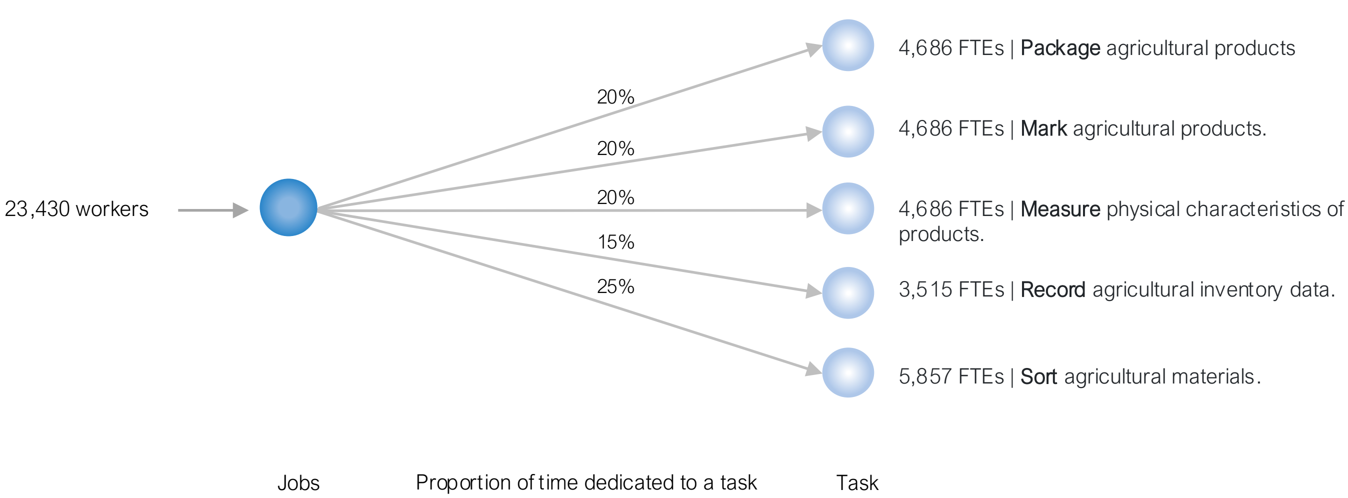

B.2 Methodology

ExampleCo – Job Profiles Crosswalk

Artificial intelligence (AI) refers to systems capable of learning from data to perform tasks that generally require human intelligence. AI encompasses a broad set of approaches, including symbolic logic, machine learning, and neural networks, and is increasingly recognized as a general-purpose technology with applications across a wide range of sectors and jobs (Eloundou et al., 2024).

We developed an AI-based crosswalk to map ExampleCo job titles to the standardized ONET taxonomy, translating employee roles into ONET's detailed task structure. This alignment is essential because ONET provides both task lists and data on task frequency for each occupation. By linking ExampleCo titles to ONET, we can allocate wages to specific tasks and conduct a task-level wage analysis. Our methodology has two stages.

1. We used Sentence-BERT (SBERT) embeddings to identify the ten most semantically similar ONET occupations for each ExampleCo title.
2. We used Sonar (Perplexity AI) to review these candidates, confirm the best match, or suggest a better alternative. This combination of embedding-based retrieval and large language model evaluation ensures comprehensive and accurate crosswalks to support task-level wage allocation.

Exhibit 5 – Illustration of the dynamic framework to map ExampleCo Workforce to O*NET job profiles.

Data Structure

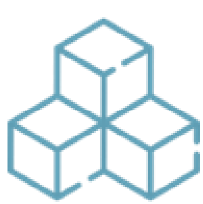

- Workforce structure (O*NET)
- ExampleCo (Employee masked info, Job descriptions)

Modeling

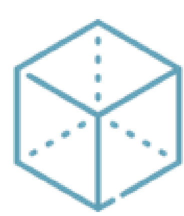

- Crosswalk with (1)BERT (2)SONAR

B.3 Methodology
Exposure of Tasks to AI

We create an “exposure metric” for each of the tasks in the job profiles we map to the ExampleCo dataset. Traditional approaches often rely on expert judgment and provide static snapshots that quickly become outdated (Brynjolfsson et al., 2018; Felten et al., 2018; Frey & Osborne, 2017; Felten et al., 2021). To address these limitations, we extended Webb (2019) by applying machine learning to extract and process information from news articles in real time, using news media as a proxy for societal attitudes toward AI adoption. To reduce bias from sensationalism and partisanship, we integrate sentiment analysis that favors neutral articles and validate our findings using two independent data sources (Liu et al., 2019; Reimers & Gurevych, 2019). Our working paper demonstrates that this approach aligns with established metrics and enables a scalable, continually updated perspective on the impact of AI on work (Bouquet & Sheffi, 2025). Within this framework, we calculate a job’s AI exposure by taking a weighted average of its tasks’ exposure, reflecting the known time distribution across tasks (Bouquet & Sheffi, 2025).

Exhibit 6 – Illustration of the dynamic framework to evaluate the exposure of a task to AI.

Modeling	Output

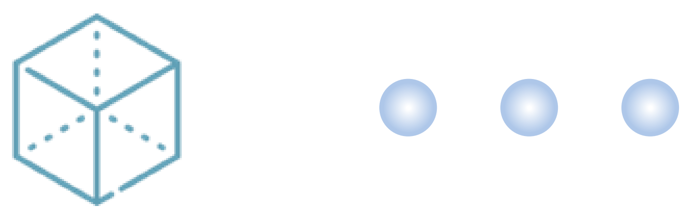

- ML pipeline for exposure metric
- Tasks exposure to AI

A practical example of the breakdown of a singular job profile (in this case an “Assistant”) can be observed in Exhibit 7. Notice how each task that constitutes the job has an associated exposure metric (High to Low, in reality these are values from 0 to 1).

Exhibit 7 – Example of a job profile broken down into eight core tasks with varying risk levels (with risk AI augmentation/automation is intended) and time spent on the task.

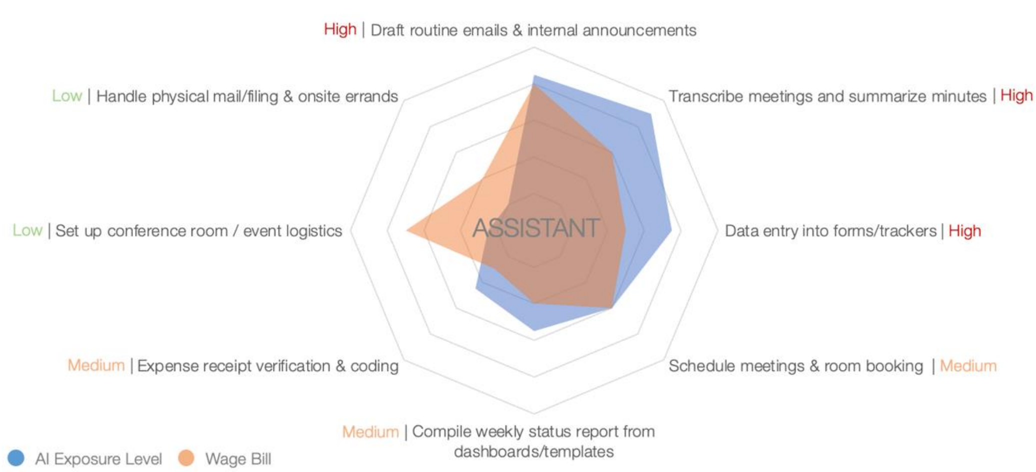

C.1 Results
Employment

How many workers in ExampleCo workforce will be impacted by AI?

2,358 (18%)
Number of ExampleCo Employees exposed to AI.

Exposed Workers

- 2,358 FTEs (~18.2% of ExampleCo's workforce) are highly exposed, meaning more than 20% in weight of there tasks are highly exposed to AI.

We assess AI exposure at both the task and job levels. Each task receives an exposure score based on a defined metric. Tasks are considered "exposed" if their score falls within the top 25 percent, creating a data-driven, percentile-based threshold. Once exposed tasks are identified, we calculate their total wage bill.

At the job level, a position is considered exposed if more than 20 percent of its tasks are AI-exposed, indicating that at least 20 percent of daily activities are susceptible to AI automation. This approach considers both the nature of tasks and their economic significance, rather than relying only on task counts or job-level averages.

Table 1 – We present a sample of tasks highly exposed to AI, along with their associated wage bill, employee count, number of departments and number of divisions performing the task.

Rank	Task Title	Wage Bill ($M)	Positions	Departments	Divisions
1	Edit documents.	2.66	531	216	53
2	Record information from meetings or other formal proceedings.	1.2	412	329	50
3	Explain project details to the general public.	1.17	1,576	369	42
4	Conduct eligibility or selection interviews.	2.26	499	143	49
5	Coordinate regulatory documentation activities.	2.03	268	48	26
6	Authorize construction activities.	7.21	1,039	154	30
7	Negotiate project specifications.	1.11	299	184	38
8	Examine financial records to ensure compliance with policies or regulations.	2.63	285	139	47
9	Edit written materials.	2.88	576	226	53
10	Review license or permit applications.	3.83	426	137	47
11	Review blueprints or specifications to determine work requirements.	1	1,443	336	30
12	Determine operational criteria or specifications.	2.86	1,444	354	42
13	Time vehicle speed or traffic-control equipment operation.	9.23	1,283	316	29
14	Survey land or bodies of water to measure or determine features.	8.57	2,338	518	34
15	Advise others on legal or regulatory compliance matters.	1.58	419	155	51

C.2 | Results
Salaries and Wage Bill

What is the impact of AI on the wage bill of ExampleCo?

$1B
Total wage bill in ExampleCo.

Total Wage bill ExampleCo

- In 2025, ExampleCo's wage bill (sum of all employee salaries) was just over one billion dollars

$181M
Total wage bill in ExampleCo exposed to AI.

Total Wage bill that is exposed to AI in ExampleCo

- This value was deduced using the 25% most exposed task based on our exposure scores. The total exposed wage bill represents 18.2% of the total ExampleCo wage bill.

Exhibit 8 – We display ExampleCo tasks with wage bills of at least $1 million.

The horizontal axis shows each task's AI augmentation potential, while the vertical axis represents total wage bill in millions. Each bubble reflects a specific task, with size indicating the number of employees and color gradient showing wage bill size. Tasks further right are more exposed to AI, and those higher on the chart have greater salary costs. For example, "edit documents" is positioned on the far right, reflecting high AI exposure. Although its wage bill is about $2.7 million, its bubble size indicates many employees perform this automatable clerical task. This example demonstrates how the chart highlights tasks that are both common and well-suited to current AI capabilities.

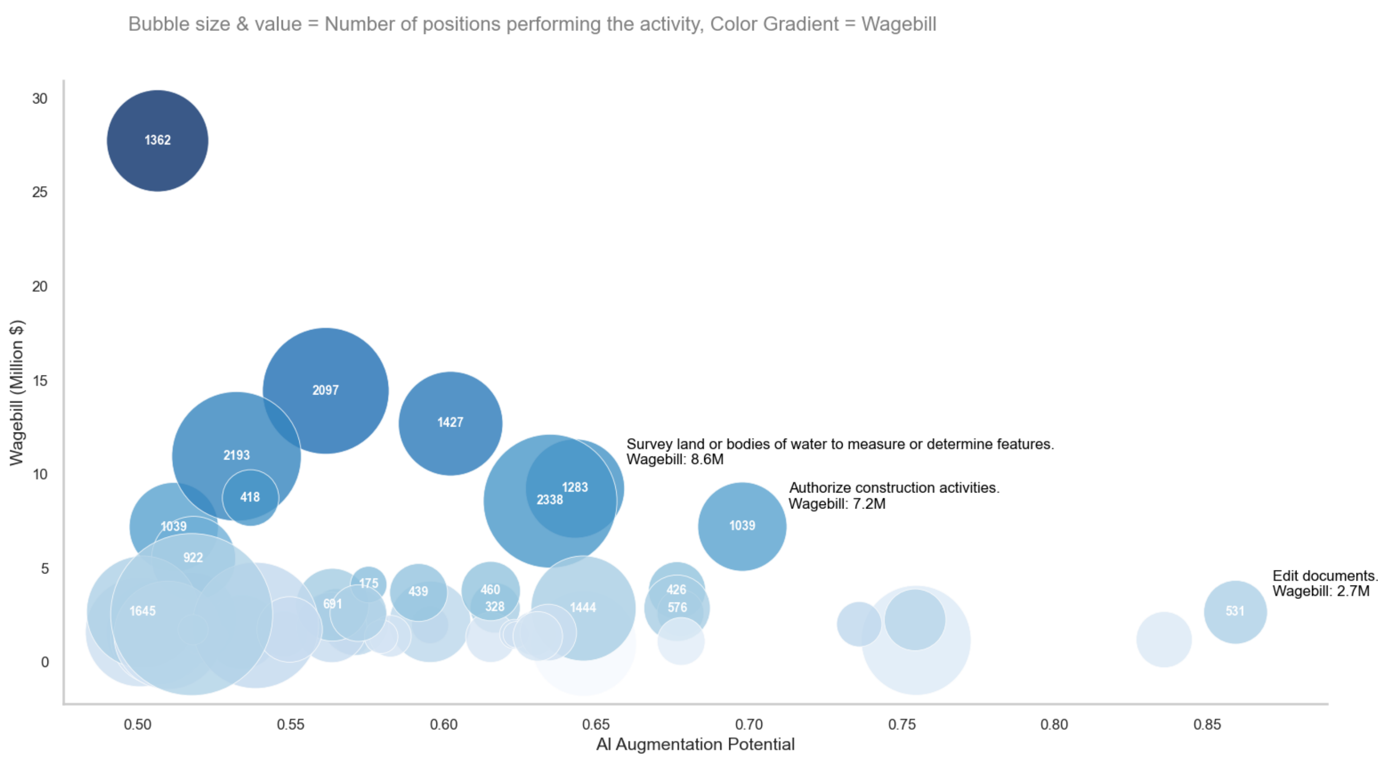

C.3 Scenarios

Simulating different deployments

What is the impact of successful AI deployment on ExampleCo?

While AI can perform tasks more efficiently, it is essential to consider the following

- AI may not completely replace tasks, but it can augment workers' productivity.
- The increased productivity from AI can free time for workers to focus on other tasks (thereby allowing fewer workers to perform certain jobs or allowing for more work to be performed).

Targeted AI Deployment And Scenario Analysis

- We evaluate several AI deployment scenarios, prioritizing tasks with high AI vulnerability scores and associated wage costs.
- Dell'Acqua et al. (2023) found that workers using AI completed tasks 25% faster and achieved higher-quality results. We assess productivity increases of 15%, 25%, and 35%.
- Each productivity scenario uses a different task selection cutoff based on exposure. Results are shown in Table 2.

For example, implementing AI for only five tasks can generate $1.4 million in agency-wide savings in our conservative scenario, and up to $2.3 million in the BCG scenario, assuming a 25% productivity increase.

Table 2 – Estimated savings under different AI-impact scenarios for selected numbers of tasks.

5 most exposed tasks with wage bill over $1M

Number of tasks	Scenario 15%	Scenario 25%	Scenario 35%
5	$1.4M	$2.3M	$3.3M

10 most exposed tasks with wage bill over $1M

Number of tasks	Scenario 15%	Scenario 25%	Scenario 35%
10	$4.0M	$6.7M	$9.4M

20 most exposed tasks with wage bill over $1M

Number of tasks	Scenario 15%	Scenario 25%	Scenario 35%
20	$8.9M	$14.8M	$20.7M

50 most exposed tasks with wage bill over $1M

Number of tasks	Scenario 15%	Scenario 25%	Scenario 35%
50	$24.3M	$40.5M	$56.7M

D.1 Case Study
"Edit documents" task

What are the impacts of implementing AI for the task "edit documents"?

$2.7M
Total wage bill of "editing documents".

Total wage bill of "editing documents" in ExampleCo

- The task is shared across 531 employees.

53
Districts in which the task "edit documents" is done.

Where is the task done?

- The task "edit documents" is done in 53 districts (95%).
- The task "edit documents" is done in 216 departments (20%).

A targeted AI solution for the "edit documents" task could significantly improve the daily work of 531 ExampleCo employees. Many hold high-responsibility, high-paying positions (see Table 3) where their primary value lies in strategic, technical, or leadership activities rather than clerical editing. Since editing is repetitive and rules-based, automating part of this workload would allow these employees to focus on higher-value tasks. This example shows how AI augmentation can deliver substantial organizational value by reducing low-leverage administrative work for skilled staff.

Table 3 – Top 10 employees who "edit documents" in ExampleCo and their associated wage bill and task cost.

PII	Job Title	Department	Division	Employees Wage Bill	AI ROI
162021	Senior Executive Advisor	Management Services	Strategic Init & Innov	$134,660	$10,495
247708	Senior Executive Advisor	Management Services	Strategic Init & Innov	$108,766	$8,477
121106	Executive Admin Assistant II	Administration	Administration	$90,226	$7,032
230422	DE/DD Executive Assistant II	Administration	Austin District	$87,046	$6,784
194904	DE/DD Executive Assistant II	Divisn Director & Staff	Prof Engineering Procure	$86,838	$6,768
145221	Executive Admin Assistant I	Administration	Administration	$85,064	$6,630
219844	DE/DD Executive Assistant II	Business Ops	Info Technology Division	$84,657	$6,598
246630	DE/DD Executive Assistant II	Administration	Beaumont District	$84,600	$6,593
148345	Commission Chief Clerk	Commission Support Office	Administration	$84,340	$164
264548	DE/DD Executive Assistant II	Administration	El Paso District	$83,892	$6,538

D.2 Case Study
Value captured from AI deployment

How can we leverage AI to reduce the wage bill and reinvest in human capital?

$675k
Estimated savings by deploying copilot.

$190k
Estimated copilot license cost.

$485k
Value captured by deploying copilot.

Savings by deploying Copilot to "edit documents"

- Based on estimated 25% efficiency increase (Dell'Acqua et al., 2023).

Annual cost of Copilot to ExampleCo

- Based on an estimated $30/month for 531 employees for the "Copilot" add-on to the standard Microsoft 365 package.

Estimated value captured by deploying Copilot

- By subtracting the cost of Copilot from the savings.

The "Edit documents" task offers a clear opportunity for targeted AI augmentation at ExampleCo. Introducing Copilot for routine document editing can free employees to focus on core activities such as stakeholder coordination, project management, regulatory work, technical reviews, and executive support. The value is direct: if AI tools can reliably accelerate or automate even part of the editing process, the resulting wage bill savings can be significant. Copilot can automate or accelerate many repetitive sub-tasks within document editing, including:

- **Document refinement and clarity**: Improving grammar, structure, and readability in drafts, memos, letters, and reports; rewriting or summarizing content such as meeting notes into concise, polished summaries.
- **Formatting and presentation:** Applying templates, correcting headings and spacing, organizing structure, and producing professional documents ready for internal or leadership review.
- **Consistency, accuracy, and compliance:** Cross-checking content against SOPs or previous documents, preparing polished email responses, and redacting or standardizing text for internal or public release.

These capabilities address clerical, repetitive work that Copilot is well suited to augment. By reducing time spent on these tasks, ExampleCo can achieve measurable wage-bill savings and redeploy employees to higher-impact work. However, realizing this value requires more than simply activating the tool. Successful AI deployment pairs technology with training. A phased rollout, such as a focused one-hour Copilot orientation on prompts, best practices, and common use cases, helps staff integrate AI into their workflow and builds confidence. This approach accelerates adoption and enables ExampleCo to capture the full benefit of targeted AI augmentation.

Copilot

D.3 Case Study
Deployment strategies

How can we strategically deploy AI in ExampleCo?”

A. ExampleCo could deploy Copilot to all 531 employees who regularly edit documents. This would provide immediate, uniform access across 53 districts and 216 departments, enabling the agency to realize productivity gains quickly. Given the rapid advancement and proven success of this technology in other industries (see introduction), delaying deployment may result in missed opportunities. However, this approach requires significant upfront licensing costs, increases training demands, and may lead to uneven adoption if some units are less prepared. A broad rollout also makes it more difficult to address department-specific needs or refine training before full implementation.

B. Alternatively, ExampleCo can start with targeted pilot projects in select departments (see Exhibit 9) before expanding statewide. This staged approach reduces initial costs and operational risks, allows for tailored training such as focused 1-hour onboarding sessions, and provides early data on productivity improvements to guide leadership decisions. Pilots also reveal operational differences, such as variations in document workflows, which help ExampleCo refine deployment and maximize value. The main drawbacks are slower realization of savings and potential perceptions of inequity among departments not included in early phases. Additionally, pilot results may not fully represent the entire organization. Nevertheless, this method offers a deliberate, data-driven path to full adoption once usage patterns and impact are validated.

Resources generated from successful project deployment can be reinvested in employee training, workflow automation, and workplace improvements. These steps enhance staff skills, streamline operations, and boost job satisfaction, helping retain talent and support future organizational growth.

Exhibit 9 – Top 10 ExampleCo departments by total “editing documents” wage bill.

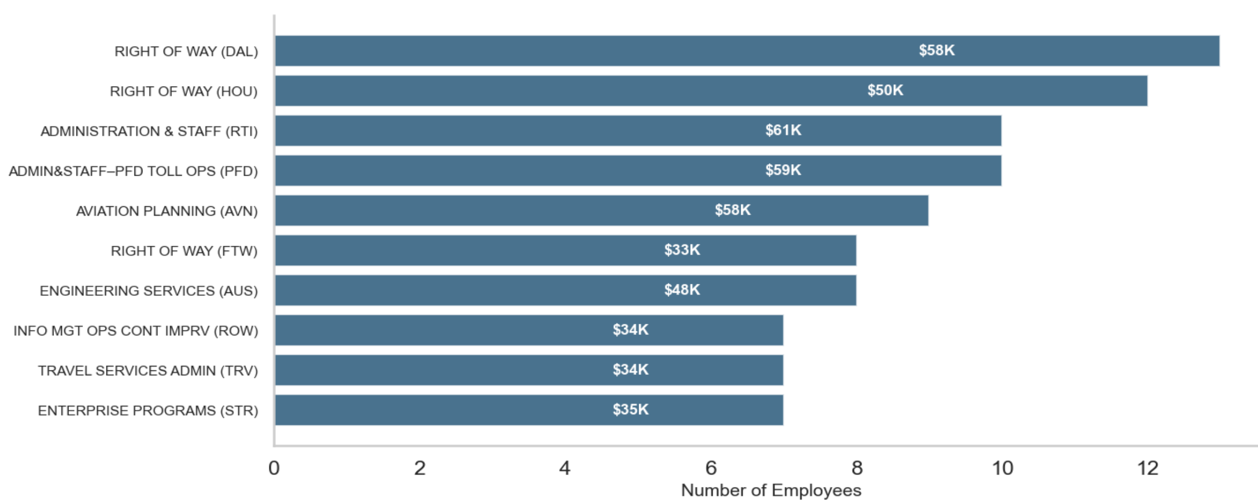

Conclusion

This report demonstrates that AI's impact on the transportation workforce is most effectively assessed at the task level rather than the job level. By breaking down ExampleCo roles into specific activities, we gain a clear and economically sound understanding of where AI can accelerate work and where operational changes may occur. Our analysis shows that approximately 18% of ExampleCo's total wage bill is associated with tasks in the top quartile of AI exposure. This proportion is financially significant and broadly distributed, making targeted, task-specific deployment the most effective strategy. The "edit documents" example highlights this point: while the task represents a modest share of total spending, it is performed by 531 employees across 53 districts and 216 departments, many in key administrative or managerial positions. Supporting even a portion of this work with AI can free up capacity for higher-value activities, such as stakeholder coordination, technical reviews, and project management. AI augmentation typically benefits multiple tasks. Tools like Microsoft Copilot also enhance related activities, including drafting correspondence, preparing summaries, cleaning documentation, and checking regulatory text. These improvements create productivity gains beyond the original task. When implemented thoughtfully, such tools can increase efficiency, reduce clerical workload, and improve work quality, while supporting reinvestment in training, operational improvements, and employee well-being. While this framework is valuable, several limitations exist. Task-level mapping relies on accurately aligning ExampleCo job descriptions with O*NET's task taxonomy. Although government roles generally align well, some mismatches or omissions are likely. Improving alignment would require larger, specialized language models and higher computational costs. As AI exposure metrics evolve rapidly, estimates must be updated regularly. AI-workforce mapping should be treated as an ongoing, iterative process rather than a one-time assessment.

Managerial Implications

- **Targeted AI Deployment:** ExampleCo now has a task-level roadmap that identifies high-cost, high-exposure tasks, enabling precise, low-risk pilot projects. Deploying Copilot to the 531 employees who "edit documents" presents an immediate opportunity. Achieving a 25% productivity gain would yield significant estimated annual savings of $485K annually, along with faster approvals, improved quality control, and better coordination across departments.
- **Spillover Effects:** Beyond financial savings from "edit documents," Copilot creates broad benefits across many administrative tasks. It assists with drafting, rewriting, summarizing, preparing minutes, standardizing templates, and quality-checking. This shifts time from clerical work to higher-value responsibilities and amplifies operational benefits.
- **Spillover Effects:** Beyond financial savings from "edit documents," Copilot creates broad benefits across many administrative tasks. It assists with drafting, rewriting, summarizing, preparing minutes, standardizing templates, and quality-checking. This shifts time from clerical work to higher-value responsibilities and amplifies operational benefits.
- **Reinvestment into People:** Resources generated from AI deployment can be reinvested in the workforce. Training and reskilling can expand administrative roles into GIS, workflow automation, or project data management. Operational tools can reduce bottlenecks, and ergonomic or flexibility improvements can strengthen retention, especially in early-tenure roles where turnover has recently increased. Reducing repetitive tasks improves both employee satisfaction and productivity.
- Organizational Resilience: Forward-looking upskilling and talent strategies can make AI a stabilizing force, preserving institutional knowledge and strengthening operational robustness across the organization.

Principal Researchers

Yossi Sheffi, PhD
Director, MIT Center for Transportation and Logistics,
Elisha Gray II Professor of Engineering Systems, Massachusetts Institute of Technology
sheffi@mit.edu

Pierre Bouquet
PhD Student, Massachusetts Institute of Technology
pibou149@mit.edu

Nicolo' Piergiovanni Bagnoli
Visiting Student, Massachusetts Institute of Technology
nicolopb@mit.edu

References

Autor, D., & Thompson, N. (2024, September 3). Schumpeter lecture: Does automation replace experts or augment expertise? The answer is yes. Massachusetts Institute of Technology.

Bouquet, P., & Sheffi, Y. (2025). News as a dynamic predictor of job automation risk (Working paper). Massachusetts Institute of Technology.

Brandes, P., & Wattenhofer, R. (2016, April 29). Opening the Frey/Osborne black box: Which tasks of a job are susceptible to computerization? arXiv preprint.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? AEA Papers and Proceedings, 108, 43–47.

Bureau of Transportation Statistics. (2024). National transportation statistics. U.S. Department of Transportation.

Dell'Acqua, F., Saran, A., McFowland, R., Krayer, L., Mollick, E., Noy, S., Brynjolfsson, E., Lifshitz-Assaf, H., Lakhani, K., & Kellogg, K. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Working paper.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), 1306–1308.

Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. AEA Papers and Proceedings, 108, 54–57.

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal, 42(12).

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280.

Hoek, R. V., DeWitt, M., Lacity, M., & Johnson, T. (2022, November 8). How Walmart automated supplier negotiations. Harvard Business Review.

Kang, J., & Choi, D. (2023). Artificial intelligence-powered innovations in the fashion industry: A mixed-methods study on AI-based customer services. International Journal of Fashion Design, Technology and Education, 17(2), 1–15.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019, July 26). RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint.

Martin, J. (2022, March 7). Developing a method for measuring time spent on green tasks. Office for National Statistics.

National Center for ONET Development. (2025). ONET OnLine. National Center for O*NET Development.

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. arXiv preprint.

Rolf, B., Jackson, I., Müller, M., Lang, S., Reggelin, T., & Ivanov, D. (2022). A review on reinforcement learning algorithms and applications in supply chain management. International Journal of Production Research, 61(20), 1–29.

Company. (2024). ExampleCo overview

Webb, M. (2019, November 6). The impact of artificial intelligence on the labor market. SSRN Working Paper.

Sheffi, Y. (2025, May 15). Managing supply chains in a tariff-fueled trade war. MIT Sloan Management Review.