

**WESTERN CREDIT UNION MANAGEMENT SCHOOL
PROJECT II
2025
(For the 2026 session)**

Amy Durst – adurst@roguecu.org

Student's Name and E-Mail Address

Rogue Credit Union

Credit Union Name

1370 Center Drive, Medford, OR 97501

Address

4-12-2026

Date Submitted

\$4.2 Billion

End of year asset size

(12/2025)

696 Spring Valley Drive, Medford, OR 97501

Student's Address

AVP of Internal Support

Credit Union Position

(270) 847-0632

Day Time Telephone Number

TABLE OF CONTENTS

I.	PROJECT I INSERTION & CURRENT CONDITIONS UPDATE	2
A.	PROJECT I INSERTION.....	2
B.	UPDATE CURRENT CONDITIONS	125
C.	RESEARCH COMPONENT.....	133
II.	PEER GROUP CONSTRUCTION & COMPARISONS	140
A.	CONSTRUCTION AND DESCRIPTION	140
B.	COMPARISONS	146
III.	STRATEGIC INFLECTION POINTS	156
IV.	SWOT ANALYSIS	164
V.	STRATEGIC GOALS & IMPLEMENTATION PLAN	177
VI.	FINANCIAL UPDATE AND THREE-YEAR FORECAST	191
A.	Three-Year Financial Forecast (2026–2028)	191
B.	7-YEAR GRAPHS.....	196
C.	FINANCIAL DATA ANALYSIS.....	201
i.	Current Year Analysis	201
ii.	Forecast Assumptions & Analysis	204
VII.	SPREAD ANALYSIS	211
VIII.	FEASIBILITY ASSESSMENT ON IMPLEMENTATION	215
IX.	EXECUTIVE SUMMARY.....	219
X.	YOUR CONCLUSIONS	223
XI.	REFERENCES	227
XII.	APPENDIX.....	230

B. UPDATE CURRENT CONDITIONS

In Project I, you completed an extensive overview of your credit union and its existing conditions.

In this section, update any **significant** changes that occurred since you completed Project I, including:

- Professional changes for you (e.g. your position, job responsibilities, credit union committee assignments, professional affiliations, etc.)
- Organizational changes within your credit union (e.g. structural changes, changes in CEO and/or other “C suite” roles, creation/termination of wholly-owned CUSOs, etc.)
- Mergers and Field of Membership expansions
- Operational and Product changes (e.g. new or closed branches, new technology implemented, new or eliminated products)

DO NOT include financial results as that information will be addressed in a different section of this project

Professional Changes:

I continue to serve as Assistant Vice President of Internal Support at Rogue Credit Union, overseeing the Deposit Operations, Internal Support, and Knowledge Management teams. While my title has remained the same since Project I, my responsibilities have continued to expand through a variety of strategic initiatives and “other duties as assigned.”

Most notably (and spoiler alert), Rogue completed a legal merger with Members 1st Credit Union in November 2025, which will be discussed more in the following sections. As part of that process, I assumed the additional role of Data Mapping Coordinator for the conversion. This role provided a unique opportunity to gain deeper insight into our core system and the complex network of third-party systems that support credit union operations. The data mapping process required coordination across multiple departments and vendors, careful tracking of thousands of data elements, and strict adherence to aggressive data-mapping and

system-cutover milestones. While Rogue's business systems analysts led the technical analysis, my role focused on maintaining forward momentum, coordinating communication across teams, managing the timeline for data mapping decisions, and ensuring issues were escalated and resolved quickly. The experience reinforced the critical importance of cross-functional collaboration during large-scale technology and merger initiatives, while also giving me a deeper appreciation for the technical expertise required to successfully execute a core conversion. Moreover, I could never secure a future as a business systems analyst.

Within my primary area of responsibility, I continue to expand Rogue's enterprise knowledge management capabilities. As the organization increasingly explores the role of artificial intelligence in supporting employees and improving service delivery, I have been invited to participate in several internal committees focused on innovation and strategy. These include the AI Steering Committee and the Digital Experience Innovation Council, where discussions center on how emerging technologies can enhance both employee productivity and member experience. Many of these discussions center on how emerging artificial intelligence tools can improve employee productivity, enhance internal knowledge access, and support faster decision-making across the organization.

I have also been asked to contribute to broader organizational initiatives through participation on the Talent & Experience Strategic Driver Steering Committee and the Financial Wellbeing Committee. These committees focus on strengthening Rogue's employee value proposition, improving team member experience, and advancing financial well-being initiatives for both employees and members. And happen to be two areas I am extremely passionate about.

In addition to my internal responsibilities, I remain actively engaged in professional and community organizations. I currently serve as the Communication Coordinator for the Southern

Oregon Chapter of Credit Unions. I am also a board member for the Southern Oregon Youth Symphony, where I serve on the organization's finance committee.

Organizational Changes:

Since the completion of Project I, Rogue Credit Union has experienced several leadership changes that reflect the organization's continued strategic evolution. In May 2025, Rogue's Chief Lending Officer, Chad Heese, departed the organization after nine years of service. His departure created an opportunity to bring in new executive leadership to guide the credit union's lending strategy during a period of continued growth and operational transformation.

In February 2026, Marcus Wertz joined Rogue Credit Union as Chief Lending Officer(CLO). Wertz brings more than 20 years of credit union leadership experience, most recently serving as Chief Member Officer and CLO at Greater Nevada Credit Union. His background includes extensive experience in consumer, mortgage, and commercial lending, centralized lending operations, credit administration, risk-based pricing, and delinquency management. In his role with Rogue, Wertz oversees senior lending leadership, including the Senior Vice President of Consumer and Mortgage Lending and the Vice President of Business Services, and is responsible for guiding enterprise lending strategy, portfolio performance, and disciplined growth.

Rogue also expanded its executive leadership structure with the creation of the Chief Experience Officer (CXO) role, now held by Scott Mulkins. Mulkins has been with Rogue for more than eleven years and previously served as Senior Vice President of Enterprise Risk and Business Intelligence. His promotion to CXO reflects the credit union's increasing focus on aligning member experience, employee experience, and financial well-being outcomes across the enterprise. In this role, Mulkins is a key contributor to Rogue's Financial Well-Being Strategy

and Operating Model, which positions financial well-being as a core organizational strategy rather than a standalone program. He also participates in cross-functional leadership forums, including the Financial Wellbeing Committee, helping integrate experience strategies with operational execution.

These leadership changes demonstrate Rogue Credit Union's commitment to strengthening executive leadership capacity while aligning strategic priorities around sustainable lending growth, member financial well-being, and enterprise-wide experience management.

Mergers and Field of Membership:

Since the completion of Project I, Rogue has experienced several developments that have expanded its geographic footprint and strengthened its growth strategy. These include the merger with Members 1st Credit Union in Northern California, the opening of a new branch in Caldwell, Idaho, and the development of a long-anticipated branch in Coos Bay, Oregon.

Rogue legally merged with Members 1st Credit Union (M1), headquartered in Redding, California, effective November 11, 2025; however, the systems merger is set to begin on June 26, 2026. Rogue serves as the surviving institution, retaining its name, charter, and headquarters in Medford, Oregon. The merger expands our presence in Northern California by adding five branch locations and approximately sixty team members. Integrating the two organizations required significant operational coordination, including core system conversion, data mapping, and regulatory reporting as systems and processes were consolidated. In my role as the merger's Data Mapping Coordinator, I have been responsible for coordinating the mapping of M1's Flex core data to our FISERV core.

At the same time, Rogue continued its branch expansion strategy in other growth markets. A new full-service branch opened in Caldwell, Idaho, strengthening the credit union's

presence in Western Idaho and supporting continued growth in retail membership, commercial lending, and mortgage services. The branch opening was widely supported by the local community and attended by city leadership, Chamber of Commerce representatives, and community organizations.

Rogue is also moving forward with the construction of a new branch in Coos Bay, Oregon. After more than a decade of planning, this new Oregon coast location will make it easier for members in the region to access services and marks a significant step in our commitment to our communities.

Together, these initiatives demonstrate an intentional approach to growth; leveraging mergers to expand into new markets while continuing to invest in physical branches that strengthen community presence and member accessibility.

Operational and Product Changes:

Over the past twelve months, Rogue has implemented a coordinated set of product launches, digital enhancements, and technology initiatives designed to deepen primary financial relationships, increase digital engagement, and improve operational efficiency. Collectively, these efforts reflect a strategic shift toward rewarding everyday banking behaviors, expanding member self-service capabilities, and allowing employees to focus more heavily on complex needs and relationship-based interactions.

Relationship-Based Deposit Products

One of the most significant product initiatives was the launch of *MyRewards Checking*, a behavior-based checking account designed to increase engagement and long-term member loyalty. The account features a tiered-rate structure with a monthly fee that can be waived when members meet participation criteria. Rewards are tied to behaviors that strengthen member

relationships, including debit card usage, electronic deposits, enrollment in paperless statements, and participation in the Save the Change program.

Members can track their progress through a dedicated dashboard within online banking, providing real-time visibility into qualification status and earned benefits. Internally, staff are supported through Member Insights tools that help identify eligible members and facilitate product conversations. Strategically, *MyRewards Checking* serves as a cornerstone relationship account, encouraging digital adoption while reinforcing behaviors that deepen member engagement.

Rogue also refreshed its savings offerings through the introduction of *MyMoney Market*, which emphasizes competitive dividend rates while maintaining liquidity and accessibility. Features such as check writing and overdraft protection linkage were retained to provide flexibility for members. The product was intentionally positioned as a complementary savings vehicle alongside *MyRewards Checking*, supporting deposit growth while providing members with a higher-earning savings option.

Payments and Card Enhancements

Rogue expanded its payments ecosystem through the implementation of Zelle, enabling real-time person-to-person payments within Rogue's digital banking environment. The rollout followed a phased approach, beginning with internal availability before expanding to broader member access. Implementing Zelle required coordination across multiple departments, including Payments, Marketing, Compliance, Training, and Internal Support, as well as updates to system logic to ensure accurate transaction counting for *MyRewards Checking* participation metrics.

Rogue also launched a new credit card product, the *Visa Signature Rewards* card, designed to align with the credit union's broader loyalty and engagement strategy. The card introduced an enhanced rewards structure intended to encourage everyday spending while strengthening Rogue's overall rewards ecosystem. Strategically, the launch supports interchange revenue growth and reinforces Rogue's value proposition as a primary financial institution for members.

Digital Self-Service Expansion

In addition to new products, Rogue significantly expanded digital self-service capabilities across its online and mobile banking platforms. These enhancements represented a deliberate shift toward digital-first delivery for routine transactions. New capabilities include expanded card controls that allow members to manage card usage digitally, real-time visibility into *MyRewards Checking* qualification status, and expanded digital account origination that enables members to open accounts and apply for products online with immediate access once approved.

Additional improvements include digital stop payment functionality, expanded internal transfer capabilities, and the transition of several routine service requests to digital forms. Collectively, these enhancements improve convenience for members while reducing operational friction and decreasing reliance on call center and branch interactions.

Artificial Intelligence Initiatives

In addition to product and digital channel enhancements, Rogue has begun exploring practical applications of artificial intelligence to support employee productivity and operational efficiency. To ensure responsible adoption, the credit union established an AI Steering Committee to evaluate proposed use cases, review security and compliance considerations, and align AI initiatives with organizational risk management practices.

Rogue has also begun providing controlled employee access to approved AI tools, including Microsoft 365 Copilot, Copilot Studio, and OpenAI's ChatGPT, paired with guidance on appropriate data usage and security guardrails. These tools are currently being used to assist with tasks such as document summarization, research, internal communications drafting, and analysis of large information sets.

To encourage hands-on experimentation, Rogue launched an internal AI Agent Builder Challenge, where employees were invited to design AI agents capable of solving real operational problems. Proposed solutions focused on improving internal knowledge access, automating routine information retrieval, and assisting staff with research and documentation. The challenge allowed the organization to identify promising use cases while providing employees with practical experience experimenting with AI in a controlled environment.

Taken together, the product launches, digital enhancements, and early AI initiatives implemented over the past year reflect Rogue's strategic focus on strengthening long-term member relationships while improving operational efficiency. Rewards and incentives were aligned with behaviors that increase engagement and reduce long-term costs, while expanded self-service capabilities allow routine transactions to migrate to digital channels. This shift enables employees to spend more time addressing complex member needs and supporting deeper financial relationships.

C. RESEARCH COMPONENT

Artificial Intelligence (AI) is predicted by some to offer the promise of transformational productivity growth for workers. (Singla et al., 2025). Others argue the promises of AI have been overstated. (Smith, 2024). Some credit unions have rushed to incorporate AI applications to better serve their members, and numerous technology companies have begun offering AI supported applications to credit unions. Other credit unions have been more reserved. But what might this technology mean for your career and role as a strategic leader for your credit union?

This year's research section seeks to support your personal development given this new technology, which will in turn support your strategic leadership at your credit union. This section is meant to be distinct from all other sections of your project. Please do not alter the other sections of your project based on your answers in this section. Further, this section is meant to be based on your own readings and insights, not from discussions with leadership within your credit union. Where do YOU think AI will affect you and your role as a credit union leader?

- 1) What is your current familiarity and usage of AI technologies to support your leadership at your credit union? Please explain with specific examples, including any tasks in your role that you already augment with AI. Has AI meaningfully changed your decision quality, speed, or strategic insight?
- 2) How, if at all, have you used AI tools to support completion of your WCMS Project? Were there specific areas where you intentionally chose to utilize, or not utilize, available AI tools to support your development as a strategic leader within your Project? Why?
- 3) Please describe and justify your personal AI usage policy as a leader within your credit union. What tasks do you (or would you) delegate to AI, what will you never delegate, and how do you remain accountable? If your organization has an AI Policy, you may reference it, but your personal AI usage policy should be personalized to you and your responsibilities.
- 4) Finally, how do you predict AI technologies will change the skills needed at your credit union in the decade ahead? Please discuss two specific skills credit union leaders should be developing now to be prepared for how AI may change credit union employment in the decade ahead. Are you currently taking steps to develop these skills?

Please present your Research Component in four distinction sections, as outlined above. Support your analysis by (correctly) citing in your answers at least two additional articles that support some of your assertions (Two articles total, not two per section. Do not count the provided articles as one of the two additional cited articles.). The two articles must be from reputable, professional publications, but do not necessarily need to be recent. The articles should be used to support or challenge your arguments and be included as part of your Project Appendix section.

1. Current Familiarity and Usage of AI

My familiarity with artificial intelligence (AI) has grown significantly over the past several years as generative AI tools have become more accessible in professional environments. While I do not have a technical background in AI development, I have become increasingly comfortable using AI as a tool to *support* research, learning, and decision-making in leadership.

In my role, which includes responsibilities related to internal support, knowledge management, and deposit operations, I frequently encounter situations that require synthesizing large amounts of information, evaluating potential solutions, and communicating complex ideas clearly. AI has proven particularly helpful in these areas. For example, I use AI tools to summarize research articles, explore alternative ways to frame strategic questions, and organize complex information into more structured outlines. And occasionally assess the tone of my writing and tell me to perhaps *not* send an email.

One of the most valuable uses of AI for me has been accelerating early-stage thinking. When approaching a complex topic, AI can help identify potential perspectives or themes that I might not have considered initially. However, I treat this output as a starting point rather than a finished product. Leadership decisions still require context, organizational knowledge, judgment, and humanity that AI systems do not possess.

Some researchers argue that AI has the potential to significantly improve productivity for knowledge workers. Singla et al. (2025), for example, suggest that generative AI tools could lead to meaningful productivity gains across many industries. Other researchers are more skeptical, noting that many claims about AI's transformational impact have historically been overstated (Smith, 2024). My experience reflects elements of both perspectives. AI can accelerate research and information synthesis, but the quality of leadership decisions still depends on human judgment and experience.

Additional research supports the idea that AI can improve performance when used to support human work rather than replace it. Brynjolfsson, Li, and Raymond (2023) found that generative AI tools

can significantly improve worker productivity on certain knowledge tasks by helping employees access relevant information and guidance more quickly. In my own experience, AI has improved the speed at which I can explore ideas and organize information, but the final interpretation and decision-making remain human responsibilities.

2. AI Usage in Completing the WCMS Project

AI tools have played a limited but intentional role in supporting the completion of my WCMS Project. I have primarily used AI to assist with research discovery, summarizing academic articles, and organizing ideas. For example, AI tools can quickly summarize lengthy research papers or help identify themes across multiple sources, allowing me to focus more attention on interpretation and analysis. I have also occasionally used AI to help organize outlines or suggest ways to structure complex sections of the project. When I first started using AI tools for writing support, I admit it felt a little like cheating. I remember when asking Microsoft Word's Clippy for help was considered cutting-edge technology, so having a tool that can reorganize ideas or summarize research almost instantly still feels a little surreal. But I also remember when "googling" something was cheating or using a calculator.

AI has also been helpful as a learning tool. When I encounter unfamiliar or particularly complex concepts, I sometimes ask the AI to explain them in several different ways until the idea clicks. It is the perfect thought partner. For example, I might ask for a technical explanation and then ask the tool to simplify it or provide a practical example. That approach has been particularly helpful when working through research-heavy material in the WCMS project.

AI has also proven useful for certain practical tasks. Do I read long contracts anymore? Absolutely not. Have I used AI to summarize one so I can understand the key points more quickly? Yes. Have I also occasionally used AI for something closer to therapy after a particularly complicated meeting? Also yes. While those examples may be lighthearted, they highlight an important point: AI can

help process information or provide perspective, but it does not replace the thinking, judgment, or accountability of a leader.

Because of that perspective, I have tried to use AI *thoughtfully* rather than heavily. The purpose of WCMS is leadership development, and outsourcing the thinking process would undermine that goal and the investment Rogue has made in my development. And honestly, Matt Stephenson, my CEO, knows exactly what we are learning at WCMS, and expects us to apply it! For this reason, I used AI primarily as a support tool for research efficiency rather than as a substitute for my own analysis or conclusions.

Listening to discussions about AI in leadership reinforced this approach. In a Harvard Business Review IdeaCast episode discussing the concept of “AI workslop,” researchers described how organizations are beginning to see an increase in low-quality, AI-generated content that appears polished but lacks meaningful insight. When leaders rely too heavily on AI to produce finished work, it can actually create additional work as employees must review, refine, and correct the output (Harvard Business Review, 2026). I do still get amused at workslop when someone starts describing better member service as “a robust, multifaceted ecosystem of transformative synergy.” That’s usually my cue that a human should step back in.

I intentionally approached AI as a tool to support thinking rather than replace it while completing my WCMS project. In practice, AI helped me move faster through the early stages of research and organization, but the analysis, conclusions, and final writing remained my responsibility. As AI tools continue to evolve, I also see value in developing a deeper understanding of how these technologies can support employees in knowledge-intensive roles. In the credit union industry, where member service and employee expertise are closely connected, leaders who learn how to thoughtfully integrate AI into daily workflows will be better positioned to support both staff effectiveness and member outcomes.

3. Personal AI Usage Policy

As artificial intelligence becomes more integrated into professional environments, leaders must develop clear personal guidelines for how these tools should be used. My personal AI usage policy is grounded in three principles: augmentation, accountability, and judgment.

First, I believe AI is most appropriate for tasks that involve information synthesis, brainstorming, drafting, and organizing complex ideas. These are areas where AI can significantly improve efficiency without replacing leadership responsibilities. For example, AI can help summarize research, generate an outline for a presentation, or suggest alternative ways to frame communication. Used appropriately, it allows leaders to move more quickly through the early stages of research and writing so they can spend more time focusing on interpretation, decision-making, and strategy.

Second, there are areas where I do not believe AI should replace human leadership responsibility. Decisions involving employees, ethical considerations, organizational culture, or strategic priorities should remain firmly within human judgment. AI can generate ideas and provide information, but it cannot fully understand the context, relationships, or long-term implications that leaders must consider when making decisions. As a result, I view AI as a support tool or thought partner, rather than a decision-maker.

Finally, accountability must remain clear. When AI is used in professional work, the leader using the tool remains fully responsible for the final outcome. I think of AI much like a highly capable research assistant: it can help gather information, organize ideas, draft communications, or suggest possible solutions, but it cannot take responsibility for the consequences of those decisions. Leaders must still verify information, apply judgment, and ensure that final decisions and communications align with the organization's values and priorities. Used in this way, AI can strengthen leadership effectiveness while preserving the accountability that effective leadership requires.

4. Future Skills Needed for Credit Union Leaders

Over the next decade, artificial intelligence is likely to reshape many aspects of credit union operations, including member service, fraud detection, lending decisions, and internal workflow management. As these technologies continue to evolve, credit union leaders will need to develop new skills to effectively guide their organizations through this transition.

I believe there are two skills in particular that will become increasingly important for credit union leaders. The first is the ability to critically evaluate AI-generated insights. AI systems can produce highly convincing responses even when the underlying information is incomplete or incorrect. Leaders must therefore develop the ability to question AI-generated outputs, verify information, and apply their own judgment before acting on recommendations. As Smith (2024) argues, one of the risks of rapid AI adoption is that organizations may overestimate the reliability of AI-generated insights without fully understanding their limitations.

The second critical skill is the ability to redesign workflows and organizational processes to effectively incorporate AI. Simply adding AI tools to existing systems will not automatically improve outcomes. Many organizations struggle to capture value from AI because they fail to redesign processes around the technology. According to McKinsey & Company (2023), the organizations that achieve the greatest benefits from AI are those that rethink how work is organized and how humans and technology interact.

For credit unions, this may involve rethinking how employees access institutional knowledge, how member questions are answered, and how internal processes are supported by technology. Leaders who understand how to thoughtfully integrate AI into employee workflows while maintaining strong human oversight will be better positioned to support both staff effectiveness and member service.

Personally, I am already taking steps to develop these skills by actively exploring how AI tools can support learning, research, and knowledge sharing in my own work. Because my role includes

responsibility for knowledge management and internal support, I see opportunities to explore how AI might help employees access information more effectively while still maintaining the human expertise and judgment that members expect from their credit union.

II. PEER GROUP CONSTRUCTION & COMPARISONS

In this section you will assemble a peer group for your credit union based on given criteria. Then you will compare the performance of your credit union with that of the selected peers. Your analysis of your peer groups will require you to recognize the significant variances, their likely causes, provide insights and draw thoughtful conclusions from the data. Your Project Reader can observe these variances as well, so it is up to you to analyze them and explain WHY they exist.

A. CONSTRUCTION AND DESCRIPTION

- Utilizing Peer-to- Peer, you are to build a comparison set of credit union data for a group of credit union peers. That **five** member “Peer Group” will include information for:
- Your credit union
- **Three** national peers that are chosen based on either total assets (+/- 30%) or membership size (+/- 20%) when compared to your credit union. Each of these 3 peers must be from a different state and may include only one credit union that is domiciled in the same state as yours.
- A National Peer Average (NPA) based upon asset size, and determined as follows:

Asset Size of Your Credit Union (as of 6/30/2025)
< \$2 million
\$2 million - \$5 million
\$5 million - \$10 million
\$10 million - \$20 million
\$20 million - \$50 million
\$50 million - \$100 million
\$100 million - \$250 million
\$250 million - \$500 million
\$500 million - \$1 billion
\$1 billion - \$10 billion
>\$10 billion

List all FOUR members of your Peer Group other than your own.

For the NPA, simply list the appropriate asset range from table above. For each of the THREE natural person credit unions, provide:

- An explanation of why you selected the credit union and which selection criteria it met (asset size or membership size)
 - A brief synopsis of information that can be obtained via direct contact with the credit union, its website, NCUA and/or other industry recognized sources.
 - ***Remember that you are REQUIRED to use appropriate American Psychological Association (APA) style footnotes and a reference list to cite third party sources.***

Discuss your observations and impressions of how your credit union is similar to the other credit union and how it differs. Some, but not all, of the potential dynamics to consider are charter structure (federal or state), share insurance provider, field of membership composition, geographic footprint served, size of the branch network, member service philosophies and strategic considerations.

Construction and Description

Utilizing the Callahan Peer-to-Peer comparison tool, a five-member peer group was assembled consisting of Rogue Credit Union, three natural-person credit union peers, and a National Peer Average (NPA). Rogue Credit Union reported approximately \$3.73 billion in assets as of June 30, 2025, placing it in the \$1 billion to \$10 billion asset band for purposes of the NPA. The three natural-person credit union peers were selected using the asset-size criterion of +/- 30%, consistent with project requirements.

The four peer group members other than Rogue Credit Union are:

- Oregon Community Credit Union (OCCU) – Oregon
- Achieva Credit Union – Florida
- Ascend Federal Credit Union – Tennessee
- National Peer Average (NPA) – Credit unions with assets from \$1 billion to \$10 billion

Although asset size was the primary screening criterion, size alone was not sufficient to build a meaningful comparison group. My initial objective was to identify peer institutions within Rogue's asset band that also operated in markets similar to Rogue's footprint, regional credit unions serving smaller communities while maintaining some adjacency to larger metropolitan areas. That proved difficult within the required asset range. As a result, I refined the peer selection approach to create a more useful comparison set: one close same-state comparator and two similarly sized institutions operating in faster-growth, metro-influenced markets.

That approach also reflects Rogue's current strategic context. Rogue continues to expand in the Treasure Valley / Boise-adjacent corridor, with branches in Caldwell, East Ontario, and New Plymouth, among other nearby communities. This makes it increasingly relevant to examine institutions operating in markets shaped by stronger population growth and broader metropolitan influence. Regional planning data estimated the 2024 population of the Boise-area region at 822,890, underscoring the scale and

continued growth of the market Rogue is entering (COMPASS, 2024). In that sense, the peer group was designed not simply to satisfy an asset-size requirement, but to support a more strategic comparison of how similarly sized credit unions perform under different demographic and economic conditions.

A secondary consideration in finalizing the peer set was whether the selected institutions demonstrated a meaningful emphasis on digital capability, data-driven member service, and emerging AI-enabled or fraud-mitigation tools. This was not the primary reason for selection, but it strengthens the relevance of the comparison to the broader themes discussed later in this project. Each of the selected peers shows evidence of investing in digital delivery, self-service access, or technology modernization that may support future analytics, fraud prevention, and member experience strategies.

Oregon Community Credit Union (OCCU)

Oregon Community Credit Union was selected as the closest comparator because it operates in the same state as Rogue and therefore provides a useful same-state benchmark within a broadly similar regulatory and economic environment. OCCU also met the required asset-size criterion, reporting assets of approximately \$3.64 billion as of the comparison date. In addition to geographic proximity, OCCU offers a valuable strategic contrast because its balance sheet is notably more loan-driven than Rogue's.

A brief review of OCCU's website and annual reporting suggests a strong emphasis on digital convenience and member-facing technology. OCCU's 2023 annual report highlights enhancements to online and mobile banking, expanded video teller hours and locations, improved online account opening, text messaging, and tap-to-pay debit cards (Oregon Community Credit Union, 2023). Those initiatives suggest a credit union that is pairing branch-based service with a strong digital push. They also, like Rogue, have recently invested heavily in Knowledge Management capabilities, which also attracted me to use them as a peer.

Rogue and OCCU are similar in several important ways. Both are Oregon-based credit unions with substantial regional footprints, a commitment to member service, and growing digital delivery capabilities. However, they also differ meaningfully. OCCU appears more aggressive in deploying assets into loans, while Rogue maintains a more liquidity-oriented balance sheet. That distinction makes OCCU especially useful as both a geographic peer and a strategic counterpoint.

Achieva Credit Union

Achieva Credit Union was selected because it met the required asset-size criterion, with assets of approximately \$3.05 billion, and because it operates in a materially different growth environment than Rogue. Headquartered in Florida, Achieva serves a market shaped by strong population migration and economic expansion. According to the U.S. Census Bureau, Florida remained one of the fastest-growing states in the country in the most recent annual population estimates, making it a relevant contrast to Rogue's more regional and historically less metro-centered footprint (U.S. Census Bureau, 2024).

Achieva was also selected because it provides insight into how a similarly sized credit union may perform in a faster-growth environment while still maintaining a relatively balanced balance sheet. Achieva's 2024 annual report notes that the credit union improved online account opening, reduced operating expenses, and added features to mitigate fraud, while also moving applications to more modern platforms (Achieva Credit Union, 2024). These efforts suggest a strategic focus on both growth and operational modernization.

Achieva is similar to Rogue in that both are large regional credit unions serving broad community-based memberships and investing in digital capabilities. However, Achieva's Florida market provides stronger demographic tailwinds that may support faster loan and membership growth. This makes Achieva a useful contrast when considering how much of performance may be driven by strategy versus market opportunity.

Ascend Federal Credit Union

Ascend Federal Credit Union was selected because it also met the required asset-size criterion, with assets of approximately \$4.58 billion, and because it operates in another high-growth, metro-influenced market. Based in Tennessee, Ascend serves a broad membership base across Middle Tennessee. Tennessee added more than 86,000 residents between 2022 and 2023, reflecting continued economic and population growth that can support increased demand for financial services (U.S. Census Bureau, 2024).

Ascend also demonstrates a significant commitment to digital access and member convenience. Its annual report highlights the availability of online video appointments, and its digital-banking ecosystem includes tools such as integrated credit score monitoring through SavvyMoney within Ascend Digital Banking (Ascend Federal Credit Union, 2025a, 2025b). These offerings suggest a strategic focus on meeting members through both physical and digital channels.

Ascend is similar to Rogue in its broad regional service orientation and emphasis on member access, but it differs in meaningful ways. It operates in a larger and faster-growing market, carries a stronger capital position, and appears to benefit from the scale advantages of a broader metropolitan footprint. Comparing Rogue to Ascend provides a useful perspective on how market growth, scale, and strategy can influence performance even among institutions of similar size.

Similarities and Differences Among the Peer Institutions

Several similarities are evident across the peer group. All four institutions are member-owned financial cooperatives operating with substantial community and regional footprints. Each maintains some combination of physical branches and digital banking capabilities, reflecting the hybrid service model that now characterizes much of the credit union industry. In addition, all four institutions show evidence of investing in technology, digital member experience, and operational modernization, even if

those efforts vary in scale and emphasis (Achieva Credit Union, 2024; Ascend Federal Credit Union, 2025a; Oregon Community Credit Union, 2023).

At the same time, the differences across the peer group are what make the comparison useful. OCCU offers the closest geographic and regulatory comparison to Rogue, but it appears to pursue a more aggressive lending posture. Achieva and Ascend operate in stronger-growth environments where population expansion and metropolitan influence likely support higher loan demand and membership growth. These market differences matter because they affect not only performance outcomes, but also strategic choices around lending, liquidity, and technology investment.

Viewed collectively, the peer institutions represent different strategic models. OCCU reflects a more loan-driven balance sheet and stronger asset deployment. Achieva represents a more balanced institution operating in a favorable growth market. Ascend combines a fast-growth market with scale and a stronger capital position. Rogue, by contrast, appears more liquidity-oriented, with a balance sheet structure that may provide greater flexibility as it continues expanding into the Treasure Valley / Boise-adjacent corridor. For that reason, this peer group provides a stronger analytical framework than a simple size-only comparison. It allows for evaluation not only of whether Rogue is performing well, but also of how its strategy compares with institutions operating under different market conditions and pursuing different approaches to growth, risk, and digital transformation. Developing a meaningful peer group required more than simply identifying institutions of similar asset size. While the Callahan Peer-to-Peer comparison tool provides numerous potential candidates within the required asset range, many differed significantly in membership structure, market characteristics, or strategic focus. As a result, constructing the final peer group became an iterative process of evaluating institutions not only for size comparability but also for similarities in community orientation, regional market dynamics, and operational complexity.

B. COMPARISONS

Evaluate your credit union's performance results against your identified peers and the NPA. Analyze the data and describe the areas where significant variances exist between your credit union and the four others in your peer group. Offer potential reasons for the variances. In doing so, factor in the dynamics cited in describing the similarities and differences between your credit union and the others.

Include all of the following elements in these peer comparisons:

- 1. Balance Sheet (\$, millions)**
 - Assets
 - Shares
 - Loans
 - Investments
 - Net Worth
- 2. Annual Growth Rates (%)**
 - Assets
 - Shares
 - Loans
 - Membership
 - Net Worth
- 3. Financial Ratios (%)**
 - Loans/Shares
 - Investments/Assets
 - Net Worth/Assets
 - Delinquency/Total Loans
 - Charged Off Loans/Total Loans
 - Efficiency, Including and Excluding Provision for Loan Losses
- 4. Spread Analysis (% , based on Average Assets)**
 - Interest income
 - Cost of Funds
 - Net Interest Margin
 - Operating Expenses
 - Provision for Loan Losses
 - Fees and Other Income
 - Non-Operating Income/Expense
 - Net Income (ROA)

Comparing Rogue Credit Union’s financial performance with OCCU, Achieva Credit Union, Ascend Federal Credit Union, and the National Peer Average (NPA) shows that similarly sized credit unions can produce very different results depending on how they deploy assets, manage risk, and position themselves within their markets. The most meaningful variances in this peer group are not simply differences in scale, but differences in balance sheet strategy. Some institutions appear to prioritize loan growth and asset utilization to maximize spread income, while others maintain more liquidity and flexibility through higher investment concentrations. Regional market conditions also matter. Institutions operating in faster-growing markets, such as Tennessee and Florida, appear to benefit from stronger tailwinds in membership and loan demand than institutions operating in slower-growth regional markets. As a result, the peer comparison is most useful when viewed not as a ranking of better or worse institutions, but as a comparison of different strategic tradeoffs involving growth, liquidity, credit risk, and profitability.

Balance Sheet Comparison

Balance Sheet - WCMS - Peer Comparisons					
Formula	Rogue (OR)	Oregon Community (OR)	Achieva (FL)	Ascend (TN)	Credit Unions \$1B-\$10B (PG)
Assets	\$3,734,398,637	\$3,638,996,883	\$3,048,657,304	\$4,579,896,504	\$2,850,726,264
Shares	\$3,184,991,020	\$3,004,309,472	\$2,770,269,976	\$3,659,642,724	\$2,406,455,948
Loans	\$2,151,338,359	\$3,283,529,841	\$2,254,038,326	\$3,354,560,517	\$2,057,173,085
Investments	\$1,295,822,004	\$238,547,315	\$603,226,581	\$976,862,649	\$640,465,247
Net Worth	\$370,813,380	\$362,195,906	\$299,608,461	\$647,881,406	\$321,359,879

This data is as of Q2 2025 and retrieved from the Callahan Peer-to-Peer comparison tool

Asset sizes across the peer group are reasonably close, ranging from approximately \$3.0 billion to \$4.6 billion, with Rogue at \$3.73 billion, OCCU at \$3.64 billion, Achieva at \$3.05 billion, and Ascend at \$4.58 billion, all above the \$2.85 billion national peer average. The more important distinction, however, is not asset size itself but how each institution allocates those assets between loans and

investments. Rogue's balance sheet differs materially from OCCU and Ascend, and those differences point to fundamentally different earnings and liquidity strategies.

OCCU and Ascend both rely more heavily on lending as a primary earnings engine. OCCU's loan-to-share ratio exceeds 100%, indicating that it has deployed nearly all available shares into loans and may require supplemental funding sources to continue growing at the same pace. That posture aligns with OCCU's emphasis on expanding lending and product offerings to support growth (Oregon Community Credit Union, 2024). Ascend also maintains a large loan portfolio relative to shares, but its balance sheet appears less stretched than OCCU's. While both institutions are loan-oriented, the similarity should not be overstated: OCCU appears to be pursuing a more aggressive, fully deployed lending posture, whereas Ascend appears to be pairing strong loan utilization with a larger capital cushion and stronger overall balance sheet stability.

Achieva sits somewhat between these two models. Its balance sheet is more balanced between loans and investments, suggesting a strategy that still emphasizes lending growth but preserves more liquidity than OCCU. That posture may be especially useful in a high-growth Florida market, where strong population growth and migration can support loan demand while still requiring balance sheet flexibility to respond to shifts in funding or pricing conditions.

Rogue stands apart because of the size of its investment portfolio. At approximately \$1.30 billion, Rogue's investments are materially higher than those of any peer in the group and far above the national peer average as a percentage of assets. This position points to more than liquidity management alone. Relative to peers, Rogue's investment portfolio appears to function as both a current earnings support and a reserve of future lending capacity. In other words, Rogue is not merely holding excess liquidity; it is maintaining optionality. The portfolio supports earnings today while preserving the ability to redirect funds into loans later if pricing, demand, or market conditions become more favorable. That

makes Rogue's balance sheet less fully deployed in the near term, but potentially more flexible over the next credit and rate cycle.

This structure may also reflect market dynamics in addition to strategy. Credit unions operating in markets where share growth outpaces loan demand often maintain higher investment balances to manage excess liquidity. The rapid rise in interest rates beginning in 2022 also made investment yields more attractive at the same time that loan demand softened in many markets. In that context, Rogue's investment concentration may represent a deliberate effort to preserve earnings while waiting for stronger or more attractive lending opportunities rather than simply a shortfall in loan production.

Capital strength is more meaningfully evaluated through net worth ratio than through net worth dollars. On that basis, Ascend stands apart with a 14.15% net worth ratio, well above the 11.27% national peer average, indicating a larger strategic capital cushion. Rogue, OCCU, and Achieva are all clustered near 10%, suggesting strong but more fully deployed capital positions. This is a more useful comparison than net worth dollars alone, since dollar net worth largely reflects institution size rather than relative capital strength. Viewed that way, Ascend's balance sheet appears not only larger, but more conservatively capitalized than the rest of the peer group.

Taken together, the balance sheet comparison highlights three distinct approaches within the peer group. OCCU appears to prioritize maximizing earning assets through aggressive loan deployment, even at the cost of tighter liquidity. Ascend also emphasizes lending, but does so from a stronger capital position and with more balance in overall structure. Achieva occupies a middle position, combining growth-market exposure with a more moderate asset mix. Rogue, by contrast, maintains a more liquidity-oriented balance sheet in which investments serve as both an earnings source and a strategic reserve. That structure may limit near-term asset utilization, but it also gives Rogue greater flexibility to respond to future opportunities and changing economic conditions.

Growth Rate Comparison

Growth Rates - WCMS - Peer Comparisons					
Peer Group	Rogue (OR)	Oregon Community (OR)	Achieva (FL)	Ascend (TN)	Credit Unions \$1B-\$10B (PG)
Assets	6.45%	8.02%	6.75%	5.23%	4.13%
Shares	4.67%	8.52%	6.32%	5.28%	5.54%
Loans	5.49%	8.92%	8.49%	8.32%	4.73%
Membership	2.13%	6.23%	5.25%	6.47%	2.58%
Net Worth	9.83%	9.89%	7.37%	7.53%	5.91%

This data is as of Q2 2025 and retrieved from the Callahan Peer-to-Peer comparison tool

Growth patterns across the peer group reflect differences in both market conditions and institutional strategy.

OCCU and Achieva report the strongest loan growth, both exceeding 8%, indicating a strong emphasis on expanding earning assets. For OCCU, this growth aligns with a more aggressive lending posture that prioritizes loan deployment and yield generation. Achieva's stronger growth likely reflects the benefits of operating in Florida, where population growth and in-migration continue to support demand for housing, consumer credit, and other financial services.

Ascend also demonstrates strong growth across multiple categories, particularly in membership, where it reports 6.47% growth, well above the national peer average of 2.58%. This likely reflects the economic and population expansion occurring in Tennessee, particularly in metro-influenced markets that support household growth, employment, and loan demand.

Rogue's growth rates are more moderate. Rogue reports asset growth of 6.45%, which exceeds the national peer average but trails OCCU. Loan growth at Rogue is 5.49%, lower than OCCU, Achieva, and Ascend, but still above the national peer average of 4.73%. These results suggest a more measured lending strategy than some peers.

Membership growth at Rogue is 2.13%, slightly below the national peer average and well below OCCU and Ascend. These differences may reflect regional population trends, competitive dynamics, or

differing approaches to growth and market expansion. However, Rogue’s net worth growth of 9.83% is one of the strongest in the peer group and nearly matches OCCU’s 9.89%, despite Rogue posting lower growth in assets, loans, shares, and membership. This suggests Rogue is converting more moderate balance sheet growth into strong capital accumulation, reflecting solid earnings performance and disciplined financial management.

Financial Ratio Comparison

Financial Ratios- WCMS - Peer Comparisons					
Formula	Rogue (OR)	Oregon Community (OR)	Achieva (FL)	Ascend (TN)	Credit Unions \$1B-\$10B (PG)
Loans/Shares	67.55%	109.29%	81.37%	91.66%	85.49%
Invs./Assets	34.70%	6.56%	19.79%	21.33%	22.47%
Net Worth/Assets	9.93%	9.95%	9.83%	14.15%	11.27%
Delinquency/Total Loans	0.79%	1.65%	0.29%	0.53%	0.76%
Charged Off Loans/Total Loans	0.97%	3.82%	0.58%	0.54%	0.76%
Efficiency Ratio (including PLL)	76.35%	80.38%	79.40%	71.49%	81.01%
Efficiency Ratio (excluding PLL)	66.54%	57.38%	73.72%	64.69%	71.56%

This data is as of Q2 2025 and retrieved from the Callahan Peer-to-Peer comparison tool

Financial ratios highlight the strategic differences in asset deployment, risk tolerance, and capital positioning across the peer group.

Loan-to-share ratios vary significantly. OCCU reports the highest ratio at 109.29%, indicating a highly loan-oriented balance sheet and an aggressive asset deployment strategy. Ascend and Achieva also maintain relatively high loan utilization levels at 91.66% and 81.37%, respectively. Rogue’s loan-to-share ratio of 67.55% is significantly lower than both the peer group and the national peer average of 85.49%, reflecting a materially different balance sheet structure.

Rogue’s lower loan penetration appears to reflect more than a conservative liquidity posture. Its 34.70% investments-to-assets ratio, compared with the national peer average of 22.47%, suggests the

investment portfolio is functioning as both a reserve of future lending capacity and a current earnings substitute. In that sense, Rogue's higher investment concentration supports liquidity and flexibility, but it also helps preserve profitability while allowing the credit union to wait for more favorable loan deployment opportunities. Rather than simply indicating excess liquidity, the portfolio appears to provide optionality.

The contrast with OCCU is especially important. OCCU's 3.82% charge-off ratio, 1.65% delinquency ratio, and 1.71% provision for loan losses suggest that its higher-yield balance sheet comes with materially greater credit cost. OCCU's results point to a more risk-tolerant lending strategy in which stronger interest income (6.47%) and net interest margin (4.53%) are achieved by deploying assets more aggressively into loans, but at the cost of significantly higher credit losses. In that sense, OCCU's profile appears to represent a clear high-risk, high-reward tradeoff.

Ascend's profile is notably different, even though it also maintains a relatively high loan-to-share ratio. Ascend combines 91.66% loan-to-share, 0.54% charge-offs, 0.53% delinquency, and the strongest ROA of 1.09% in the peer group. This suggests Ascend is pairing strong loan utilization with better overall credit performance and profitability, making its strategy meaningfully different from OCCU's more credit-cost-intensive model.

Capital strength is more meaningfully assessed through net worth ratio than net worth dollars. On that basis, Ascend stands apart with a 14.15% net worth ratio, well above the national peer average of 11.27%, suggesting a larger capital cushion and a more conservative capital posture. Rogue, OCCU, and Achieva are all clustered near 10%, indicating strong but more fully deployed capital positions. This provides a more accurate comparison than looking at net worth dollars alone, which largely reflects size.

Credit quality metrics remain generally healthy across most of the peer group, though OCCU is a clear outlier. Achieva and Ascend both maintain lower delinquency and charge-off ratios than Rogue

and OCCU, suggesting either a more favorable portfolio mix, stronger underwriting performance, or better recent credit conditions in their markets.

Efficiency ratios also reveal a deeper story than the headline numbers suggest. OCCU reports the strongest efficiency ratio excluding provision for loan losses at 57.38%, indicating a very efficient revenue model before credit costs are considered. However, when provision expense is included, OCCU's efficiency ratio rises sharply to 80.38%, which is weaker than Rogue's 76.35% and well above Ascend's 71.49%. This indicates that much of OCCU's pre-provision efficiency advantage is offset by the cost of supporting its more aggressive lending strategy. Rogue's efficiency ratio is somewhat higher than the strongest performers, but still better than the national peer average, indicating generally effective expense management.

Spread Analysis

Spread Analysis- WCMS - Peer Comparisons					
Formula	Rogue (OR)	Oregon Community (OR)	Achieva (FL)	Ascend (TN)	Credit Unions \$1B-\$10B (PG)
Interest Income	5.63%	6.47%	5.00%	5.01%	5.03%
Cost of Funds	1.73%	1.94%	1.50%	2.08%	1.92%
Net Interest Margin	3.90%	4.53%	3.50%	2.93%	3.11%
Operating Exps.	3.31%	3.43%	3.46%	2.71%	3.00%
P.L.L.	0.64%	1.71%	0.34%	0.40%	0.49%
Fees & Other Inc.	1.08%	1.45%	1.20%	1.26%	1.08%
Non-Op. Inc./Exp.	0.01%	0.00%	0.03%	0.01%	0.05%
Net Inc. (R.O.A.)	1.03%	0.84%	0.93%	1.09%	0.75%

This data is as of Q2 2025 and retrieved from the Callahan Peer-to-Peer comparison tool

Institutions with higher loan concentrations generally generate stronger interest income, and that pattern is evident within the peer group. OCCU reports interest income equal to 6.47% of average assets, the highest within the group, while Rogue's 5.63% exceeds the national peer average but remains below

OCCU. This reflects the fact that OCCU is deploying a larger share of assets into loans, which typically generate higher yields than investments.

Cost of funds across the peer group is relatively consistent, ranging from 1.50% to 2.08%, suggesting similar funding pressures overall. Rogue's 1.73% cost of funds is slightly below the national peer average, indicating effective deposit pricing and funding management. These factors produce different net interest margins. OCCU reports the highest margin at 4.53%, reflecting its highly loan-driven structure. Rogue's 3.90% margin exceeds the national peer average and Achieva, but trails OCCU. However, margin alone does not tell the full story. OCCU's provision expense of 1.71% materially reduces the benefit of its higher margin, while Rogue's lower provision expense allows more of its spread to flow through to earnings. This highlights the tradeoff between gross yield and credit cost.

Operating expenses vary modestly across the group. Rogue reports operating expenses of 3.31% of average assets, slightly above the national peer average but comparable to peers. Ascend's lower operating expense ratio of 2.71% may reflect scale advantages or a more efficient operating model.

Provision for loan losses vary more significantly and help explain differences in final profitability. OCCU's elevated provision aligns with its higher charge-offs and delinquency, while Rogue's provision remains moderate relative to its balance sheet structure. Ascend's lower provision, combined with high loan utilization, helps explain why it produces the strongest ROA at 1.09%.

Despite different balance sheet strategies, most institutions in the peer group generate returns on assets near or above 1%, indicating generally strong profitability. Rogue's 1.03% ROA remains competitive, particularly given its more liquidity-oriented structure.

Summary: Insights Learned

The peer comparison demonstrates that credit unions of similar size can pursue very different strategic approaches to growth, earnings, and risk management. OCCU appears to be pursuing a more

aggressive, yield-maximizing lending strategy. Its high loan-to-share ratio and strong margin support earnings, but those benefits come with materially higher delinquency, charge-offs, provision expense, and weaker post-provision efficiency. Achieva appears to follow a more balanced model in a favorable Florida growth market, benefiting from strong demographic tailwinds without showing the same degree of credit stress. Ascend also operates with relatively high loan utilization, but its results suggest a different and more balanced strategy than OCCU's. Ascend combines strong loan deployment with lower credit losses, better overall efficiency, a stronger capital ratio, and the highest ROA in the peer group. In that sense, Ascend appears to be pairing growth with stronger risk control.

Rogue stands apart for its significantly higher investment concentration and lower loan-to-share ratio. That structure appears to reflect more than liquidity management alone. Rogue's investment portfolio serves as both a current earnings support and a reserve of future lending capacity, allowing the credit union to preserve flexibility while remaining profitable. Although Rogue is less aggressive in asset deployment than OCCU or Ascend, it continues to generate strong net worth growth and competitive profitability. Overall, the peer comparison suggests that Rogue's strategy emphasizes optionality, liquidity, and capital accumulation, while peers such as OCCU and Ascend place greater emphasis on loan-driven earnings. The national peer average generally falls between these approaches, underscoring that multiple balance sheet strategies can support strong financial performance, though with different tradeoffs in risk, flexibility, and return.

III. STRATEGIC INFLECTION POINTS

An inflection point is a significant external event or trend that changes the course of an entire industry, sector, or economy. These forces are often outside the control of any single organization and can bring both risk and opportunity. Understanding and preparing for such shifts is a vital leadership competency.

Inflection points are:

- Broad in scope
- External to your credit union
- Universal across the credit union system
- They are not issues that are specific to your institution, region, or field of membership.

There are many such outside influences impacting your credit union that should be considered beyond looking at only peer or other industry data. Gaining the perspectives of your credit union's leadership on these inflection points can be a key determinant in its future strategies and success.

This differs from a SWOT analysis, as SWOT analysis evaluates both internal (strengths and weaknesses) and external (opportunities and threats) factors that are specific to your credit union's current position. You will be asked to complete a SWOT in the next section.

In contrast, this assignment is focused only on **broad, external trends** that all credit unions are navigating, regardless of size, charter, or geography.

- Select 3–5 external inflection points currently influencing the credit union industry.
- You may choose from the list below or identify your own, provided they meet the criteria above
 - Changing competitive landscape including non-financial companies
 - Rapid adoption/expectation on the use of technology for services
 - Baby boomers retiring
 - Millennials entering their peak earning years
 - Changes in the employment base – job skills, automation, etc.
 - Shift to subscription-based economy (e.g. automobiles)
 - Evolving consumer expectations relative to products, services and delivery channels
- Schedule a strategic conversation with your CEO and/or members of the Senior Management Team and ask for their perspectives on how these trends are impacting the credit union and how the leadership of your credit union is thinking about the risk and opportunity.
- Using the information gleaned from these discussions and your own insights, write an assessment that discusses the impacts of the items you selected. Also include discussion on how your credit union plans to mitigate the risks and/or capitalize on the potential opportunities presented by these inflection points.

Strategic Inflection Points

Credit unions operate within a financial services environment shaped by external forces that no single institution can control. These strategic inflection points influence how financial institutions compete, deliver services, and respond to evolving member expectations. While individual credit unions cannot control these forces, understanding them allows leadership to proactively adapt strategies, allocate resources effectively, and position the organization for long-term success.

Conversations with Rogue Credit Union's senior leadership reinforced that several broad external trends are currently reshaping the credit union industry. (Pickens, 2026; Stephenson, 2026). These trends include rapid technological advancement, changing consumer expectations, evolving demographics, and growing competition from non-traditional financial providers. Collectively, these forces are altering how members interact with financial institutions and how financial institutions must deliver value.

One of the most significant implications of these trends is the increasing need to balance technological innovation with personalized member service. As digital capabilities expand, credit unions must ensure that technology enhances, not replaces, the human relationships that distinguish the cooperative financial model. Artificial intelligence, in particular, has the potential to improve operational efficiency while empowering employees to provide more personalized financial guidance and support. Research suggests that financial institutions that successfully integrate AI into their workflows can redesign processes, improve service delivery, and create greater value for both employees and customers (McKinsey & Company, 2025).

The following five inflection points represent several of the most significant external forces currently shaping the credit union industry: artificial intelligence, evolving consumer expectations, fragmentation of financial relationships, demographic shifts, and the expansion of non-traditional competitors into financial services.

#1 Artificial Intelligence

Artificial intelligence (AI) is rapidly transforming the financial services industry by enabling organizations to automate processes, analyze large volumes of data, and enhance service delivery. Financial institutions are increasingly investing in AI-driven tools to improve operational efficiency, personalize services, and reduce friction in customer interactions. McKinsey & Company estimates that generative AI and advanced analytics could add \$200 billion to \$340 billion in annual value to the global banking sector, largely through productivity improvements and enhanced customer experiences (McKinsey & Company, 2023).

Across financial services, AI applications are already being used to support fraud detection, automate customer service interactions, analyze transaction data, and assist employees in providing accurate information to customers. A recent Deloitte survey found that more than 70% of financial institutions are increasing their investment in artificial intelligence and automation as part of broader digital transformation initiatives (Deloitte, 2024). Similarly, research from Ron Shelvin at Cornerstone Advisors indicates that financial institutions increasingly view AI, automation, and advanced analytics as key priorities for improving operational efficiency and customer experience (Shevlin, 2024).

While AI is often discussed primarily in terms of automation and cost reduction, its strategic importance extends beyond operational efficiency. When implemented thoughtfully, AI can enhance employee effectiveness by reducing repetitive administrative tasks, surfacing relevant information quickly, and supporting more informed decision-making. Credit union research suggests that artificial intelligence has the potential to improve both operational productivity and member experience by equipping employees with better information and decision-support tools (Filene Research Institute, 2023).

Rogue's leadership has recognized artificial intelligence as a strategic enabler rather than simply an operational tool. Strategic planning discussions emphasize the responsible use of AI to enhance

member experiences, improve operational scalability, and strengthen employee effectiveness while maintaining strong governance and data security standards. Rogue has already begun integrating AI-enabled knowledge management and workflow tools designed to provide employees with real-time access to accurate information when assisting members (Rogue Credit Union, 2025).

By reducing the time employees spend searching for policies, procedures, and operational guidance, these tools allow staff to focus more attention on financial guidance, problem resolution, and relationship building. In this way, artificial intelligence can act as a force multiplier for relationship-based service by freeing employees to concentrate on higher-value member interactions. As adoption of AI continues to expand across financial services, credit unions will need to develop strategies that balance innovation, operational efficiency, and governance while ensuring that technology strengthens the trust-based service model that defines the cooperative financial system.

#2 Evolving Member Expectations

Member expectations regarding financial services have changed significantly in recent years. Today's consumers increasingly expect financial services to be fast, convenient, and accessible through digital channels. These expectations are shaped not only by other financial institutions but also by technology companies that have set new standards for seamless digital experiences across industries. Research indicates that digital banking has become the dominant channel for many consumers. According to the American Bankers Association, 71% of consumers now prefer managing their bank accounts through mobile or online channels, while fewer than 10% identify physical branches as their primary banking channel (American Bankers Association, 2023).

For credit unions, evolving consumer expectations create both operational and strategic challenges. Members increasingly expect digital self-service capabilities, real-time access to information, and consistent service experiences across channels. Institutions that fail to meet these

expectations risk losing engagement and loyalty, particularly among younger members who are accustomed to mobile-first experiences.

At the same time, financial decisions such as purchasing a home, managing debt, or planning for retirement still require personalized guidance and trust. This creates an opportunity for credit unions to differentiate themselves by combining digital convenience with human expertise.

Our leadership has identified digital self-service as an important response to these changing expectations. Strategic planning materials emphasize providing members with tools to independently access information and complete transactions while ensuring that employees remain available to deliver personalized financial guidance when needed. Investments in internal workflow improvements and employee enablement further reflect an understanding that member experience is shaped by the entire service ecosystem. When employees have access to accurate information and efficient processes, they are better positioned to provide timely and consistent support to members.

As member expectations continue to evolve, credit unions must design service models that integrate digital convenience with personalized human guidance, ensuring that technology enhances member relationships rather than reducing financial interactions to purely transactional exchanges.

#3 Fragmentation of Financial Relationships

Another significant trend affecting the credit union industry is the fragmentation of financial relationships. Historically, consumers often maintained most of their financial activity with a single financial institution. Today, however, many individuals distribute their financial activity across multiple providers. Digital wallets, fintech applications, and specialized financial platforms have contributed to this shift by offering convenient solutions for specific financial tasks such as payments, budgeting, or peer-to-peer transfers. According to the Consumer Financial Protection Bureau (CFPB), more than three-quarters of U.S. consumers now use at least one payment app. (CFPB, 2023).

This fragmentation presents a strategic challenge for credit unions because it reduces the likelihood that any single institution will remain central in a member's financial life. When financial activity is distributed across multiple providers, institutions may experience lower deposit balances, reduced transaction volume, and fewer opportunities to deepen long-term financial relationships.

Andrew Lepczyk's article suggests that strengthening core relationships, particularly through share draft accounts and lending relationships, remains important to increasing primary financial institution status and expanding member engagement across products (Lepczyk, 2024). Rogue's leadership has similarly recognized that maintaining primary financial institution status requires deeper engagement across multiple products and services. Strategic planning materials emphasize participation-based loyalty strategies designed to encourage members to consolidate financial activity within the cooperative model. Programs such as Rogue Rewards reinforce the value of maintaining multiple relationships with the credit union. Leadership discussions also highlight financial guidance and trust as key defenses against fragmentation. By positioning Rogue as a long-term partner in members' financial well-being rather than simply a transactional provider, the credit union seeks to remain relevant even as financial activity becomes increasingly distributed across multiple platforms.

#4 Demographic Shifts

Demographic changes represent another important inflection point for the credit union industry. The retirement of the baby boomer generation and the financial behaviors of younger generations are reshaping financial services demand and delivery models. In the United States, approximately 10,000 baby boomers reach retirement age each day, creating increased demand for financial services related to retirement planning, income management, and long-term financial security (Pew Research Center, 2023).

At the same time, younger generations are entering the financial system with different expectations regarding technology, convenience, and financial guidance. Economic conditions have also

influenced financial behaviors across generations. Rising housing costs have increased the median age of first-time homebuyers to 36 years old, the highest level on record (National Association of Realtors, 2024).

These demographic changes require financial institutions to support members across multiple life stages while adapting service delivery to meet different expectations. Younger members often prioritize digital accessibility and convenience, while older members may place greater emphasis on trust, financial advice, and personal relationships.

Rogue's strategic planning explicitly recognizes the importance of serving diverse demographic segments, including youth, working families, Spanish-speaking communities, and pre-retirees. These groups represent key growth opportunities within Rogue's field of membership. By investing in both digital capabilities and employee training, we aim to meet the expectations of younger members while maintaining the human connection that distinguishes the credit union model. These demographic shifts reinforce the need for service models that combine digital accessibility with personalized financial guidance throughout the member lifecycle.

#5 Non-Traditional Competitors

The entry of non-traditional competitors into financial services represents another significant inflection point for the credit union industry. Technology companies, fintech startups, and digital payment providers are increasingly offering financial products and services that compete directly with traditional financial institutions. Companies such as Apple, PayPal, and Square have embedded financial services within broader technology ecosystems, enabling consumers to make payments, access credit, and manage financial transactions without relying exclusively on traditional financial institutions. According to Accenture, more than 60% of consumers globally now use fintech services for payments, transfers, or financial management (Accenture, 2023).

These companies are reshaping consumer expectations by delivering intuitive digital experiences, real-time transactions, and seamless integration with everyday digital platforms. For credit unions, this evolving competitive landscape presents both risks and opportunities. While fintech providers may capture individual components of the financial relationship, they often lack the trust, community presence, and long-term relationship orientation that characterize credit unions.

Rogue recognizes that fintech firms and technology companies increasingly compete for specific elements of the financial relationship. Strategic discussions emphasize that our competitive advantage lies in combining modern digital capabilities with cooperative ownership, community engagement, and personalized service. Rather than attempting to replicate fintech business models, Rogue's strategy focuses on leveraging trust, financial guidance, and personalized relationships while selectively adopting technology and strategic partnerships that enhance member experience.

Summary

Taken together, these inflection points illustrate that the credit union industry is experiencing significant transformation driven by technological innovation, evolving member expectations, demographic change, and intensifying competition. These forces are reshaping how financial institutions deliver services and how members evaluate financial relationships.

For Rogue Credit Union and all credit unions, these trends reinforce the importance of combining digital capability with relationship-based service. Artificial intelligence and emerging technologies offer opportunities to improve operational efficiency and service consistency, but their greatest value lies in enabling employees to focus on meaningful member interactions. Credit unions that successfully leverage technology to enhance employee capability while preserving trust and personal service will be best positioned to thrive in this evolving financial landscape.

IV. SWOT ANALYSIS

SWOT analysis is a strategic tool used to assess an organization's position by examining both internal and external factors. The acronym SWOT stands for:

- Strengths: internal attributes or resources that support success
- Weaknesses: internal limitations or areas in need of improvement
- Opportunities: external conditions the credit union could use to its advantage
- Threats: external conditions that could challenge or undermine success

When applied effectively, SWOT analysis helps an organization align its capabilities with its environment in order to define a clear, sustainable strategy.

Earlier in this project, you completed a peer financial comparison and explored strategic inflection points—broad, industry-wide trends likely to impact all credit unions. This next step asks you to build on that work by identifying specific internal and external factors unique to your own credit union.

It's important to distinguish between the external factors used in a SWOT analysis (Opportunities and Threats) and the inflection points previously discussed. While both sets of factors are external in nature, inflection points refer to macro-level industry forces that apply broadly across the credit union system. In contrast, the opportunities and threats in a SWOT analysis are typically more localized and specific to your field of membership, regional dynamics, or competitive context. For this reason, please avoid reusing inflection points as SWOT items, and instead select a distinct set of external factors for this portion of the project.

To complete your SWOT analysis, identify 3-5 items in each SWOT category—strengths, weaknesses, opportunities, and threats—resulting in a total of 12 to 20 items. At least two items in each category must address areas of the credit union outside your direct area of responsibility. (If you are the President or CEO, this requirement does not apply.)

You are expected to write a comprehensive and fully developed analysis. For each item in your SWOT, provide a detailed and thoughtful discussion that explains why the item was selected and how it relates to your credit union's strategic position. Support each entry with specific context, observations, or data when possible. Your analysis should demonstrate critical thinking and reflect your understanding of both the internal dynamics and external environment of your organization. This section will directly inform the development of your strategic goals, business plan, and financial projections later in the project, so it should be approached with the same level of depth and care.

The following SWOT analysis examines Rogue Credit Union’s internal strengths and weaknesses, as well as the opportunities and threats posed by external developments affecting Rogue’s markets and strategic position. This analysis reflects the organization’s current strategic position, considering its service model, growth ambitions, competitive environment, and evolving technological landscape.

SWOT - ROGUE CREDIT UNION	
STRENGTHES	WEAKNESSES
1. Strong Relationship-Based Service Culture	1. Complexity of Technology Adoption
2. Clear Enterprise Strategic Alignment	2. Limited Scale Relative to Larger Competitors
3. Knowledge Management & Employee Enablement	3. Capacity Strain Resulting During Expansion
4. Participation-Driven Member Relationship Model	4. Uneven Enterprise Process Maturity
5. Strong Financial Capacity for Strategic Investment	5. Limited Formal Product Management Discipline in Deposits
OPPORTUNITIES	THREATS
1. Deeper Household Penetration Within Existing Membership	1. Intensifying Competition in Rogue’s Regional Markets
2. Growth in Underserved Regional Demographic Segments	2. Deposit Pressure from Alternative Savings and Cash Management Options
3. Personalized Engagement Within Local Markets	3. Rising Service Expectations Across Channels
4. Selective Partnerships That Extend Capability	4. Competition for Specialized Talent in a Constrained Labor Market
5. Expanded Advisory Role in Members’ Financial Lives	5. Integration Risk Associated with Future Growth and Merger Activity

Strengths

1. Strong Relationship-Based Service Culture

One of Rogue Credit Union’s most significant strengths is its long-standing commitment to relationship-based service and member advocacy. Enterprise strategy consistently emphasizes *Personal Delivery Experience*, defined as delivering exceptional service in the moments that matter most to members. This approach prioritizes trust, human connection, and financial guidance rather than competing solely on transactional convenience or pricing.

Relationship-based service remains a meaningful differentiator for credit unions. Research indicates that trust plays a critical role in consumer financial decision-making, particularly when individuals are navigating complex financial choices such as lending, retirement planning, and debt management. The Edelman Trust Barometer consistently finds that financial institutions perceived as trustworthy are significantly more likely to retain customer loyalty and long-term engagement (Edelman, 2024). Institutions that cultivate strong personal relationships with their members are therefore better positioned to maintain engagement even as digital banking options expand.

For Rogue, this cultural foundation is reinforced by the work of the Digital Experience team, which intentionally designs digital tools to support, rather than replace, relationship-based service. Through the use of member personas and journey mapping grounded in UX research and behavioral insight, Rogue ensures its digital experiences reflect the same empathy, trust, and clarity that define in-person interactions. These personas help the organization anticipate member needs, reduce friction in moments of uncertainty, and create seamless transitions between digital self-service and human guidance when it matters most. As financial services become increasingly digitized and fragmented, Rogue's ability to align persona-driven digital design with its Personal Delivery Experience strengthens its role as a trusted financial partner rather than simply a transactional service provider.

2. Clear Enterprise Strategic Alignment

Rogue benefits from a clearly articulated enterprise strategic framework supported by well-defined strategic drivers and priorities. Key enterprise drivers, including Digital Delivery Experience, Personal Delivery Experience, Financial Soundness, and Talent and Employee Experience, provide a consistent lens for leadership decision-making and investment prioritization. Strategic clarity improves organizational execution by ensuring that initiatives across departments support common enterprise goals. Technology investments, process improvements, and service enhancements are therefore evaluated not as isolated projects but as components of an integrated strategy.

Strategic alignment improves organizational effectiveness by reducing fragmentation and ensuring that technology investments, operational improvements, and service initiatives reinforce one another. This clarity becomes especially important during periods of rapid change, such as technology modernization, market expansion, or merger integration, when coordinated execution is essential for maintaining operational stability and strategic focus.

3. Knowledge Management and Employee Enablement

Rogue has strengthened its knowledge management and employee enablement capabilities by investing in systems that improve how employees access and apply institutional knowledge. As financial services operations become increasingly complex, frontline and support employees must be able to quickly locate accurate policies, procedures, and product information in order to provide reliable member support.

Knowledge management plays an important role in enabling employees to deliver consistent service across multiple channels. Rather than relying on informal information sharing or individual expertise, knowledge management systems allow organizations to capture institutional knowledge and make it accessible to employees in real time.

One example of this capability is Rogue's development of Ask Ollie⁶, an internal knowledge platform that provides centralized access to operational guidance, procedures, and policy information. By digitizing institutional knowledge and making it easily searchable, Rogue enables employees to respond to member questions more efficiently while maintaining consistency across branches, contact centers, and digital channels. The implementation of Ask Ollie and the establishment of a dedicated Knowledge Management team have strengthened Rogue's ability to capture and distribute institutional

⁶ Ask Ollie is Rogue Credit Union's centralized internal knowledge management platform by eGain, designed to provide employees with timely, accurate, and consistent information to support service delivery and operational decision-making.

knowledge. Since launch, 100% of Rogue’s member-facing employees use the tool at least once per month, with average usage ranging from five to eight searches per employee per day.

By strengthening its knowledge management infrastructure, Rogue ensures that employees across branches, contact centers, and support teams have access to accurate and up-to-date information. This improves service consistency, reduces operational errors, and enables faster resolution of member inquiries. Over time, these capabilities also create a foundation for advanced analytics, automation, and artificial intelligence applications that can further improve operational efficiency.

4. Participation-Driven Member Relationship Model

Rogue’s strategy emphasizes participation through core financial relationships such as checking accounts, debit card usage, lending activity, and digital engagement. This participation-driven model encourages members to view Rogue as their primary financial institution rather than one of several financial service providers. Products such as *MyRewards Checking*, *MyMoney Market*, Zelle integration, and enhanced card rewards are intentionally designed to encourage everyday usage and primary financial institution behavior.

Participation is strategically important because it increases the depth and stability of member relationships. Members who rely on an institution for multiple financial services tend to demonstrate higher loyalty and lower attrition rates. They are also more likely to engage in additional services over time, such as lending products, investment services, or financial planning support.

This model also strengthens financial performance. Institutions with deeper member relationships often experience higher lifetime member value, stronger cross-product adoption, and more stable deposit funding. By designing products and services that encourage everyday engagement, Rogue strengthens both member relationships and organizational resilience.

5. Strong Financial Capacity for Strategic Investment

Rogue's financial strength provides an important foundation for long-term strategic investment. Maintaining strong capital ratios, sound balance sheet management, and disciplined risk oversight allows the organization to invest in technology modernization, talent development, and service enhancements without compromising financial stability.

Financial capacity is particularly important in the current financial services environment, where institutions face increasing pressure to invest in digital platforms, cybersecurity infrastructure, and data analytics capabilities. These investments often require high upfront costs before producing measurable returns.

Organizations with strong financial resources are better positioned to absorb these investments while maintaining operational stability. Rogue's financial capacity allows leadership to pursue innovation and growth initiatives while maintaining the safety and soundness required of regulated financial institutions.

Weaknesses

1. Complexity of Technology Adoption

As Rogue expands its use of digital tools, automation, and AI-enabled capabilities, implementation complexity increases. Initiatives such as expanded digital self-service, real-time payments, AI-supported knowledge management (e.g., Ask Ollie), and workflow automation require coordinated changes across core systems, delivery channels, policies, governance models, and employee training.

This weakness matters because the value of technology is not created by implementation alone but by effective organizational absorption. When rollout velocity outpaces readiness, adoption can be uneven across branches and support teams, and expected efficiencies may be delayed. Rogue's challenge is therefore not simply selecting the right technologies, but integrating them consistently across a growing, multi-state enterprise while maintaining service quality during transition.

2. Limited Scale Relative to Larger Competitors

Although Rogue is a strong and financially sound regional credit union, it does not operate at the scale of national banks, large fintech firms, or technology-first financial platforms. Larger competitors benefit from broader capital bases, deeper pools of specialized talent, greater product development capacity, and more extensive data and analytics infrastructure.

This relative scale disadvantage affects speed and strategic optionality. Rogue must be more selective about where it invests internal resources, when it partners with third-party providers, and how it sequences initiatives. While this does not prevent Rogue from competing effectively, it requires disciplined prioritization and a clear focus on areas where relationship-based service, participation, and trust provide differentiation rather than attempting to match scale-driven competitors' feature-for-feature.

3. Current Capacity Strain During Expansion

Rogue's growth strategy, including merger activity, field-of-membership expansion, and enterprise transformation, places additional strain on systems and teams. Recent and ongoing integration efforts require extensive cross-functional coordination, policy harmonization, data and product mapping, process redesign, and elevated levels of internal support for frontline employees.

This weakness is important because growth does not automatically translate into operational simplicity. During periods of expansion, organizational bandwidth tightens, and leadership attention may shift toward integration work, reducing capacity for other strategic initiatives. If not actively managed, these dynamics can result in short-term inefficiencies, increased employee workload, and temporary service friction even as Rogue strengthens its long-term market position.

4. Uneven Enterprise Process Maturity

While Rogue has developed strong capabilities in areas such as internal support, knowledge management, and select digital workflows, process maturity is not yet consistent across all functions.

Some areas operate with higher levels of standardization, documentation, automation, and governance than others. This matters because inconsistency can undermine both employee effectiveness and member experience. When teams operate with varying levels of process clarity and maturity, the enterprise may struggle to deliver the seamless, connected experience envisioned in Rogue’s strategic framework. Continued focus on enterprise process design, standardization, and shared operating principles will be necessary to support scalability and reduce reliance on informal workarounds as complexity increases.

5. Limited Formal Product Management Discipline in Deposits

As Rogue’s product set expands and competition intensifies, the lack of a consistently defined product management discipline *within the deposits area* represents a strategic weakness. Ownership for deposit product performance, lifecycle management, pricing strategy, and ongoing enhancement is not always clearly centralized, which can diffuse accountability across teams.

This matters because deposit products are increasingly shaped by experience design, digital functionality, pricing strategy, and ongoing lifecycle management. A more formalized deposit product management approach would strengthen Rogue’s ability to monitor performance, respond to market signals, refine participation-driven checking and related digital services, and ensure the deposit strategy remains aligned with broader enterprise goals and member expectations.

Opportunities

1. Deeper Household Penetration Within Existing Membership

One of Rogue’s most significant external opportunities is to deepen relationships within its existing membership and household base. Rather than relying solely on new-member acquisition, Rogue has an opportunity to increase participation across checking, deposits, lending, payments, and advisory relationships among members it already serves.

This opportunity is specific to Rogue’s relationship-based service model and participation strategy. Products such as participation-driven checking, complementary savings offerings, and integrated digital services provide a foundation for increasing wallet share within existing households. By leveraging trust, personalized service, and participation incentives, Rogue can become more central to members’ day-to-day financial lives, driving stronger retention, higher product engagement, and improved long-term profitability within its current field of membership.

2. Growth in Underserved Regional Demographic Segments

Rogue has a meaningful opportunity to expand its relevance and market share by serving demographic segments within its footprint that remain underpenetrated or underserved. These include younger households, Spanish-speaking communities, and members transitioning into retirement, groups that are well represented across Rogue’s markets in Southern Oregon, Northern California, and Idaho.

By tailoring outreach, financial education, product positioning, and service delivery to the specific needs of these segments, rather than relying on broad national approaches, Rogue can earn and strengthen trust, increase participation, and expand household relationships within communities it is already positioned to serve.

3. Personalized Engagement Within Local Markets

Rogue has an opportunity to enhance member engagement by delivering more personalized experiences tailored to the financial behaviors and life stages of members within its local markets. Relationship depth, transaction history, and participation patterns provide insights that can help Rogue anticipate member needs and deliver more relevant financial guidance. When combined with Rogue’s relationship-based service model, personalized engagement can strengthen financial well-being conversations, improve product alignment, and deepen long-term relationships while reinforcing trust.

4. Selective Partnerships That Extend Capability

Rogue has an opportunity to accelerate strategic priorities through selective partnerships with fintech firms, vendors, and solution providers. Partnerships can support capabilities such as digital onboarding, automation, payments, analytics, and internal enablement without requiring Rogue to build every solution internally. By selectively partnering while maintaining governance and control over the member relationship, Rogue can extend its capabilities in areas where speed and technical depth matter, while preserving its cooperative values and relationship-based differentiation.

5. Expanded Advisory Role in Members' Financial Lives

Economic uncertainty, retirement transitions, rising household debt, and increasing financial complexity create an opportunity for Rogue to expand its role as a trusted financial guide. Members increasingly seek guidance with financial planning, decision-making, and long-term financial health rather than purely transactional services. By strengthening financial wellness initiatives, education, and advisory conversations across branches, digital channels, and employee-enabled interactions, Rogue can deepen trust, reinforce loyalty, and position itself as a long-term partner in members' financial well-being.

Threats

1. Intensifying Competition in Rogue's Regional Markets

Rogue faces increasing competitive pressure from community banks, regional banks, national financial institutions, and digital-first providers operating within or adjacent to its core markets in Southern Oregon, Northern California, and Western Idaho. Heightened competition can influence loan pricing, deposit retention, and brand visibility within Rogue's established footprint. Even with strong member trust and market presence, competitors with larger marketing budgets, faster product

development cycles, or more aggressive promotional strategies may erode growth if Rogue does not continue to clearly differentiate on experience, participation, and relationship depth.

2. Deposit Pressure from Alternative Savings and Cash Management Options

A growing threat to Rogue is the increasing availability of alternative savings vehicles, including online banks, brokerage cash management accounts, and high-yield digital accounts offered by both fintechs and large financial institutions. Members may maintain transactional or lending relationships with Rogue while shifting savings balances to external providers offering higher yields or perceived convenience. This behavior can weaken deposit funding stability and reduce overall relationship depth, increasing reliance on pricing incentives to retain balances.

3. Rising Service Expectations Across Channels

Members increasingly evaluate financial institutions based on the best service experiences they encounter across industries. As expectations for digital convenience, responsiveness, and seamless interactions increase, even minor service friction can become more visible to members. For Rogue, the challenge lies in maintaining consistent service quality across branches, digital self-service platforms, and internal support functions. Any visible gaps between member expectations and actual service experiences may erode trust and loyalty if not actively managed.

4. Competition for Specialized Talent in a Constrained Labor Market

Rogue faces an ongoing threat in attracting and retaining talent with skills in analytics, digital delivery, product management, technology, and change leadership. This is especially true given that our locations are outside large metropolitan areas such as Portland and Sacramento. Competition for the mentioned capabilities is particularly intense given Rogue's regional footprint and the ability of larger

institutions and technology firms to offer higher compensation, remote work options, or narrower specialization. Because many of Rogue’s strategic priorities depend on specialized skills, the inability to recruit and retain these capabilities could slow execution of digital initiatives, data-driven engagement strategies, and service innovation.

5. Integration Risk Associated with Future Growth and Merger Activity

As Rogue continues to pursue growth through mergers and market expansion, future integrations may introduce operational complexity across systems, policies, member experience, and organizational culture. While mergers can create opportunities to expand Rogue’s geographic footprint and strengthen its competitive position, they also introduce risks that extend beyond operational integration.

One of the most significant risks associated with continued growth is the potential dilution of Rogue’s relationship-based service culture. Credit union leaders have emphasized that preserving Rogue’s culture and service philosophy during periods of expansion is essential to maintaining the organization’s long-term differentiation. During leadership interviews conducted as part of this research, Matt Stephenson and Jeanne Pickens noted that protecting Rogue’s culture and member-focused values during merger activity is a persistent concern and a factor that requires careful attention during integration planning (Pickens, 2026; Stephenson, 2026). If cultural alignment is not successfully maintained, service consistency and employee engagement could be affected, weakening one of Rogue’s most important competitive strengths.

In addition to cultural considerations, each merger introduces unique challenges related to technology integration, policy alignment, and operational coordination. Poorly paced or poorly sequenced growth could strain organizational capacity or introduce service disruptions that affect both employees and members. Successfully managing integration risk will therefore be essential to ensuring

that future growth strengthens Rogue's strategic position while preserving the relationship-based service model that differentiates the organization

Summary

Collectively, these strengths, weaknesses, opportunities, and threats illustrate a strategic environment in which Rogue Credit Union possesses strong relational and financial capabilities but must carefully manage operational complexity and rising competitive pressure. These dynamics inform the strategic priorities and goals outlined in the following section.

V. STRATEGIC GOALS & IMPLEMENTATION PLAN

You will now define five measurable goals that your credit union should pursue over the next three years. These goals must be strategically significant, forward-looking, and based on your earlier analysis of inflection points and your SWOT assessment.

At least three of your goals must address enterprise-level issues, challenges or opportunities that impact the organization as a whole. You may include up to two goals that draw on your own area of expertise or departmental insight.

Important: Goals must represent new strategic initiatives not already in progress or listed in your credit union's current strategic plan. This is your opportunity to propose meaningful, original strategies based on your learning through this project.

Each goal should be written using the SMART format (Specific, Measurable, Achievable, Relevant, Time-bound). You will use the template below to explain the goal, outline the required steps, and describe its anticipated impact. The template only provides three steps, but your implementation plan should be comprehensive and include all steps necessary to accomplish the goal.

Where applicable, estimate not only the cost of each action step but also any revenue-generating or cost-saving impact it may have. This demonstrates your understanding of the financial return and strategic value of each initiative.” This is a common oversight—your implementation plan must reflect realistic resource considerations.

STRATEGIC GOAL TEMPLATE

(Complete this for each goal. A sample is provided on the next page.)

Goal Title: [Insert a clear, descriptive title]

Goal Statement: [Write a SMART-formatted goal statement]

Goal Explanation and Strategic Rationale: [Explain why this goal matters strategically, including its connection to your SWOT or inflection point analysis]

Action Steps:

1. [Insert action step]
 - Cost/Revenue:
 - Completed by (date):
2. [Insert action step]
 - Cost/Revenue:
 - Completed by (date):
3. [Insert action step]
 - Cost/Revenue:
 - Completed by (date):

Expected Impact/Results: [Describe the strategic benefit and measurable outcome, including overall financial impacts expected if this goal is achieved].

Strategic Goal 1

Goal Title: Enterprise AI-Enabled Operational Efficiency

Goal Statement: Implement AI-enabled automation, workflow redesign, and decision-support tools across key operational and service functions by December 31, 2028, reducing manual processing time by 25% and improving Rogue’s efficiency ratio to below 50%, while maintaining service quality and member satisfaction.

Goal Explanation and Strategic Rationale: The financial services industry is experiencing rapid technological change as artificial intelligence, workflow automation, and digital tools reshape how organizations deliver service and manage operational workloads. These changes represent a strategic inflection point for credit unions as they seek to balance rising operating costs with members’ growing expectations for speed, accuracy, and convenience.

My SWOT analysis highlights both the opportunity and the challenge associated with this shift. While Rogue has strong capabilities in knowledge management and internal support systems, continued growth, merger integration, and expanding service complexity place increasing demands on operational teams. Without targeted efficiency improvements, operating costs may continue to rise faster than desired and place pressure on long-term profitability.

Implementing AI-enabled operational tools provides Rogue with an opportunity to improve efficiency while strengthening its relationship-based service model. By automating routine tasks, improving internal information access, and streamlining decision workflows, employees can spend less time performing manual processing and more time assisting members with complex financial needs. This aligns with Rogue’s broader strategic priorities of delivering exceptional member experiences while maintaining financial soundness.

From a financial perspective, this initiative is expected to require more meaningful upfront investment in 2026 than originally anticipated, particularly in technology implementation, workflow

redesign, integration, and change management. However, those near-term costs are expected to support gradual improvements in operating performance over the full forecast period. Rather than producing immediate dramatic gains, this initiative is expected to create a transition year in 2026, followed by stronger realization in 2027 and 2028.

Action Steps

1. Identify and prioritize operational processes suitable for AI-enabled automation

Conduct an enterprise-wide assessment of operational workflows to identify high-volume processes with the greatest potential for automation, decision-support improvements, and workflow redesign.

- **Cost/Revenue:** \$100,000 for process analysis, consulting support, and implementation planning. Identification of automation opportunities expected to reduce manual processing time and improve staff productivity.
- **Completed by:** December 31, 2026

2. Implement AI-enabled workflow automation and decision-support tools within high-volume operational functions

Deploy automation and AI-supported decision tools in targeted operational areas such as internal support services, operational servicing workflows, and knowledge-based employee support.

- **Cost/Revenue:** \$1,200,000 for technology implementation, vendor support, process redesign, and system integration. Expected reduction in operational workload equivalent to approximately 8–10 full-time employee capacity, enabling productivity gains without proportional staffing growth.
- **Completed by:** December 31, 2026

3. Integrate AI-powered knowledge and decision tools into frontline service channels and operational workflows

Expand the use of AI-enabled knowledge systems to support frontline employees in delivering faster and more consistent member service across branches, contact centers, and digital channels.

- **Cost/Revenue:** \$700,000 for system enhancements, service-channel integration, data connectivity, employee training, and change management. Reduced internal information search time and faster resolution of member inquiries across service channels.
- **Completed by:** December 31, 2027

4. Establish governance and performance monitoring for AI-enabled operational tools

Develop governance frameworks and performance metrics to monitor adoption, operational efficiency improvements, and service quality associated with AI-enabled systems.

- **Cost/Revenue:** Internal resources. Sustained operational efficiency, stronger governance, and consistent adoption of AI-enabled tools across the organization.
- **Completed by:** December 31, 2028

Expected Impact and Results:

Successful implementation of AI-enabled operational tools is expected to improve employee productivity, reduce manual processing workloads, and moderate the growth rate of operating expenses over time. The forecast assumes 2026 will function as a heavier transition and implementation year, with expense pressure remaining visible before benefits are more fully realized in 2027 and 2028.

Over the three-year period, these investments are expected to support gradual improvements in Rogue's efficiency ratio, bringing it below 50% by 2028, while maintaining strong service quality and member satisfaction. Operational efficiencies generated through automation will allow Rogue to redirect employee capacity toward higher-value activities such as member financial guidance, relationship development, and more complex service support.

Strategic Goal 2

Goal Title: Strengthen Primary Financial Institution Relationships Through Member Participation

Goal Statement: Increase Rogue Credit Union's role as members' primary financial institution by expanding participation across checking, lending, savings, and digital payment services by December 31, 2028, increasing the percentage of households with three or more active products by at least 15%.

Goal Explanation and Strategic Rationale:

Deepening member relationships represents one of Rogue Credit Union's most significant strategic opportunities. Many members maintain financial relationships with multiple institutions, using different providers for checking, lending, and payment services. Expanding participation across multiple financial products strengthens member loyalty while increasing the long-term value of each member relationship.

Rogue's SWOT analysis highlights the strength of its relationship-based service culture and participation-driven strategy. Products such as MyRewards Checking and integrated digital payment services encourage members to use Rogue for everyday financial activities. Increasing participation across products strengthens Rogue's role as members' primary financial institution while improving the stability of deposits and lending relationships.

From a financial perspective, deeper participation increases household engagement, improves deposit retention, and drives sustainable loan growth while reinforcing Rogue's mission of helping members achieve financial success.

Action Steps

1. Expand participation-driven product incentives and engagement programs

Develop enhanced rewards and incentives that encourage members to use Rogue checking, payments, and lending services as their primary financial platform.

- **Cost/Revenue:** \$120,000 marketing and program development costs. Increased product adoption and stronger household participation across services.

- **Completed by:** December 31, 2026

2. Enhance digital engagement tools that encourage everyday financial participation

Improve mobile and online banking features that support routine financial activities such as payments, transfers, and account monitoring.

- **Cost/Revenue:** \$150,000 digital platform enhancements. Increased digital payment usage and improved member engagement.
- **Completed by:** December 31, 2026

3. Strengthen frontline employee training focused on relationship development

Provide targeted training that equips employees to identify member needs and guide members toward broader financial participation.

- **Cost/Revenue:** \$40,000 training investment. Improved cross-product adoption and stronger member loyalty.
- **Completed by:** December 31, 2026

Expected Impact/Results:

Increasing participation across checking, lending, savings, and digital payment services is expected to deepen member relationships and strengthen Rogue Credit Union's position as members' primary financial institution. Members who maintain multiple product relationships typically demonstrate higher loyalty, increased service usage, and greater long-term engagement with the credit union.

Expanding participation will also support Rogue's long-term financial performance. As members adopt additional products and services, the credit union benefits from stronger loan demand, improved deposit stability, and increased utilization of everyday financial services such as debit transactions and digital payments.

From a financial perspective, these initiatives support projected **loan growth of approximately 7–8% annually** and **deposit growth of approximately 5–6% annually** during the forecast period.

Strengthening participation across member households also improves long-term relationship value and reinforces Rogue’s mission of helping members achieve financial well-being.

Strategic Goal 3

Goal Title: Expand Digital and Data Capabilities to Support Personalized Member Experiences

Goal Statement: Enhance Rogue Credit Union’s digital and data capabilities by December 31, 2028, to increase digital self-service transactions by 25% and support more personalized member engagement.

Goal Explanation and Strategic Rationale:

Member expectations for financial services continue to evolve as digital technologies reshape how consumers interact with financial institutions. Members increasingly expect convenient digital services that provide seamless access to financial information and services while still maintaining the ability to receive personalized support. By strengthening digital platforms and leveraging data insights, Rogue can deliver more personalized member experiences while improving operational efficiency and service accessibility.

This goal addresses the inflection point of rising digital member expectations while also leveraging Rogue’s opportunity to use data and technology to deepen relationships and improve service accessibility.

Action Steps

1. Expand digital self-service capabilities for routine member transactions

Enhance digital banking platforms to allow members to complete routine transactions quickly and independently.

- **Cost/Revenue:** \$300,000 technology investment and system integration. Reduced service workload and improved member convenience.
- **Completed by:** December 31, 2026

2. Implement enhanced member analytics and segmentation tools

Develop data analytics capabilities that allow Rogue to better understand member behavior and financial needs.

- **Cost/Revenue:** \$150,000 analytics platform enhancements. Improved marketing effectiveness and targeted member engagement.
- **Completed by:** December 31, 2026

3. Integrate personalized insights into member communication channels

Use member data insights to deliver more relevant financial guidance and product recommendations.

- **Cost/Revenue:** \$100,000 system integration and communication platform improvements. Increased member engagement and cross-product adoption.
- **Completed by:** December 31, 2026

Expected Impact/Results:

Expanding digital and data capabilities will improve Rogue's ability to deliver convenient, personalized financial services while maintaining the relationship-based service model that differentiates the credit union. Enhanced digital platforms will allow members to complete routine transactions quickly and efficiently, while improved data insights will support more proactive and personalized financial guidance.

Increased digital self-service adoption is expected to reduce the volume of routine transactions handled through higher-cost service channels such as branches and contact centers. This shift allows employees to dedicate more time to complex member needs, financial guidance, and relationship development.

Financially, increased digital adoption will help moderate operating expense growth while supporting stronger member engagement and participation across Rogue's financial products and services. Over time, improved digital capabilities also position the credit union to respond more

effectively to evolving member expectations and competitive pressures within the financial services industry. While infrastructure and capability investments occur earlier in the forecast period, the strongest operating and engagement benefits are expected to emerge more fully in 2027 and 2028 as adoption increases.

Strategic Goal 4

Goal Title: Strengthen Enterprise Process Governance and Deposit Product Ownership

Goal Statement:

Establish formal enterprise deposit product ownership and process governance structures by December 31, 2026, to improve operational consistency, accountability, pricing discipline, and product performance management across Rogue Credit Union.

Goal Explanation and Strategic Rationale:

As Rogue continues to grow, operational complexity across products, systems, and internal workflows also increases. This challenge is especially important in deposit products, where pricing, participation strategy, member experience, and operational execution must work together in a coordinated way. My SWOT analysis identified growing organizational complexity, process maturity needs, and the opportunity to strengthen participation-driven relationships through more intentional product management. Establishing clearer deposit product ownership and stronger enterprise process governance directly addresses those needs.

This goal is strategically important because deposit products are not only service offerings; they are also foundational drivers of member participation, funding stability, and long-term financial performance. Without clear ownership and governance, product decisions can become fragmented across departments, slowing execution and limiting Rogue's ability to respond consistently to changing member needs and market conditions.

Formalizing product ownership and process governance will improve accountability for product performance, support more disciplined pricing decisions, and strengthen coordination across operations, service, marketing, and analytics. It will also help Rogue scale more effectively by reducing process variability and improving consistency as the organization grows.

Action Steps

1. Define enterprise deposit product ownership roles and governance structures

Establish clear leadership accountability for deposit and participation product performance, including responsibility for pricing, performance monitoring, member engagement, and cross-functional coordination.

- **Cost/Revenue:** Internal leadership resources. Improved strategic accountability, faster product decision-making, and stronger coordination across functions.
- **Completed by:** June 30, 2026

2. Develop standardized deposit product performance dashboards and reporting tools

Create centralized reporting tools that track product performance, pricing effectiveness, participation trends, and member behavior across key deposit products.

- **Cost/Revenue:** \$50,000 analytics and reporting tools. Create centralized reporting tools that track product performance, pricing effectiveness, participation trends, and member behavior across key deposit products.
- **Completed by:** December 31, 2026

3. Implement enterprise process documentation and improvement frameworks

Standardize operational workflows associated with deposit products across departments to improve consistency, training, scalability, and execution quality.

- **Cost/Revenue:** \$75,000 for process improvement initiatives. Improved operational efficiency, reduced service variability, and fewer disruptions tied to growth or system change.

- **Completed by:** December 31, 2026

Expected Impact/Results:

Strengthening enterprise process governance and deposit product ownership will improve operational consistency and create clearer accountability for product performance across the organization. Defined ownership structures will allow leadership teams to monitor deposit product performance more effectively, coordinate improvements across departments, and respond more quickly to changing market conditions and member behavior.

This goal also supports Rogue's long-term scalability. Standardized workflows and stronger documentation reduce operational variability, improve training consistency, and minimize service disruptions associated with growth, product change, or system updates. Just as importantly, stronger governance helps ensure that deposit products are managed intentionally rather than operationally inherited.

From a financial perspective, improved process governance and product ownership are expected to support more effective pricing discipline, stronger deposit product performance, and better funding management over time. This should contribute to healthier deposit growth, improved participation across core relationship products, and better control of funding costs as Rogue continues to scale. In the forecast, that discipline supports the broader objective of maintaining solid deposit growth while gradually improving overall financial performance.

Strategic Goal 5

Goal Title: Maintain Financial Strength While Funding Strategic Innovation

Goal Statement: Maintain Rogue Credit Union's strong financial position by sustaining a net worth ratio above 10% and improving return on assets from approximately 0.85% in 2026 to near 1.00% by December 31, 2028, while continuing to invest in technology modernization and member experience improvements.

Goal Explanation and Strategic Rationale:

Financial strength provides the foundation for long-term strategic investment and organizational resilience. Rogue’s capital position enables the credit union to invest in technology modernization, operational improvements, and member service initiatives while maintaining overall financial stability.

Maintaining disciplined financial management is especially important because the forecast assumes 2026 will function as a transition and investment year. Rogue is expected to absorb higher near-term costs associated with automation, digital enablement, member onboarding, and process redesign before stronger operating benefits are realized in 2027 and 2028. This means the credit union must manage capital, earnings, and growth carefully so that strategic investment does not weaken financial soundness.

This goal reflects the need to balance innovation with discipline. Rogue’s strategy is not to maximize short-term earnings at the expense of long-term capability. Instead, the credit union is expected to preserve strong capital, accept some temporary earnings pressure in the near term, and improve profitability over time as strategic investments begin to produce stronger operating leverage.

Action Steps

1. Maintain disciplined balance sheet and capital management practices

Monitor asset growth, loan portfolio performance, liquidity, and capital levels to preserve long-term financial stability during the forecast period.

- **Cost/Revenue:** Internal financial management resources. Sustained capital strength, sound liquidity management, and continued financial resilience.
- **Completed by:** Ongoing through December 31, 2028

2. Prioritize strategic investments in digital infrastructure, automation, and technology modernization.

Allocate resources toward technology improvements that enhance service delivery, operational

efficiency, and long-term scalability, while sequencing those investments in a way that preserves overall financial strength.

- **Cost/Revenue:** \$500,000 multi-year technology investment, in addition to larger 2026 implementation and integration expenditures reflected in Goal 1 and Goal 3. Improved operational capabilities, stronger member experience, and longer-term productivity benefits.
- **Completed by:** December 31, 2028

3. Conduct regular strategic and financial performance reviews

Evaluate financial performance and progress toward strategic objectives to ensure continued alignment between investment activity, balance sheet performance, and long-term profitability expectations.

- **Cost/Revenue:** Internal leadership oversight. Improved financial discipline, stronger strategic alignment, and earlier course correction when needed.
- **Completed by:** Ongoing through December 31, 2028

Expected Impact/Results

Maintaining strong financial performance while investing in strategic capabilities ensures Rogue Credit Union remains resilient in a rapidly evolving financial services environment. Sustaining capital above 10% provides the financial flexibility needed to support technology modernization, operational improvements, and member experience initiatives without undermining the credit union's stability.

A strong financial position also allows Rogue to respond effectively to economic uncertainty, competitive pressures, and regulatory changes. By maintaining disciplined financial management practices, the credit union can continue investing in innovation while preserving long-term strength.

From a financial perspective, Rogue is expected to remain well-capitalized throughout the forecast period, with the net worth ratio staying above 10%. Return on assets is expected to soften in 2026 as Rogue absorbs higher transition-year investment, then improve in 2027 and 2028 as the benefits of automation, digital enablement, and stronger member participation are more fully realized. This

pattern supports a believable financial story of investment first, stronger realization second, and fuller payoff third.

Goals Alignment Summary

Collectively, these strategic goals align Rogue Credit Union's operational priorities with the financial drivers that shape long-term performance. Goal 1 and Goal 3 place greater pressure on 2026 expenses as Rogue invests more heavily in automation, digital capability, and service redesign, while Goal 2 supports the stronger loan and deposit growth assumptions embedded in the forecast. Goal 4 strengthens the governance and product discipline needed to sustain those gains, and Goal 5 provides the financial guardrails that keep the model aligned with capital and profitability expectations. Together, these goals create a believable sequence of investment first, stronger realization second, and fuller payoff third.

VI. FINANCIAL UPDATE AND THREE-YEAR FORECAST

Incorporating Balance Sheet, Income Statement, Key Data & Ratio Information

In Project I, you analyzed specific elements of your credit union's financial performance using mid-year data, which was used for consistency and expediency across all participants. In this section, you will now work with full year-end financial data to develop a more complete and forward-looking analysis.

You will begin by compiling your credit union's actual year-end results for 2022–2025, and then use those historical figures as the foundation to develop your own three-year financial forecast for 2026–2028. All historical data should be sourced from Callahan's Peer-to-Peer tool.

Important: The 2025 year-end data will not be available until late February 2026. Since your final project is due approximately six weeks later, it is critical that you complete the collection and analysis of your 2022–2024 year-end data before the end of 2025. Delaying this step may place you at risk of falling behind during the most time-sensitive phase of the project.

A. Three-Year Financial Forecast (2026–2028)

This forecast represents your strategic financial projection of what is likely to occur over the next three years. It should reflect your understanding of industry trends, historical performance, and your credit union's strategic direction. While precise accuracy is not expected, your forecast should demonstrate alignment with your earlier analyses and thoughtful assumptions based on the work you have already completed. You will be expected to detail your rationale for your forecasts later on in this section.

Your financial forecast must be informed by:

- The external trends and inflection points identified earlier
- The internal factors and gaps revealed in your SWOT analysis
- The financial impact of the five strategic goals you developed in the previous section

Key Factors to Consider

When developing your forecast, consider how your credit union's growth strategy will affect both income and expenses:

Income Considerations:

- Loan growth and product mix
- Investment income projections
- Fee income and other non-interest sources

Expense Considerations:

- Changes in deposit strategy or funding mix
- Staffing levels and compensation planning
- Branch or facility expansion
- Marketing initiatives
- Technology and infrastructure investments

You may also incorporate macroeconomic factors if you believe they are likely to influence your forecast. This could include expected changes in interest rates, inflation, recessionary trends, or other external shocks.

Reminder: You must incorporate the financial impact of each of your five strategic goals into this forecast. These should include both revenue-generating and cost-incurring aspects, as appropriate. Your ability to connect strategic priorities to financial outcomes is a key part of this section.

An Excel-based projection tool has been provided to support students who may be less familiar with forecasting. Use of this tool is optional, and it will not be submitted as part of your project. It is offered as a resource to guide your thinking and help organize your assumptions.

Please complete your financial forecast using the formatted tables provided on the following pages. These tables are structured to capture both historical data (2022–2025) and forecasted results (2026–2028).

Balance Sheet

Balance Sheet - Rogue Credit Union							
Formula	Dec-2022	Dec-2023	Dec-2024	Dec-2025	Dec-2026	Dec-2027	Dec-2028
Loans							
Auto Loans	\$775,473,252	\$714,397,567	\$786,864,738	\$991,341,956	\$1,070,649,312	\$1,161,654,501	\$1,266,203,409
Real Estate Loans	\$709,791,272	\$791,078,105	\$875,699,020	\$1,015,503,629	\$1,076,433,847	\$1,151,784,216	\$1,235,288,572
Credit Card & Unsecured Loans	\$140,005,897	\$158,619,475	\$143,116,874	\$138,944,632	\$152,839,095	\$168,123,004	\$185,775,919
Other Loans	\$281,847,604	\$292,078,115	\$333,464,830	\$352,878,214	\$373,652,967	\$392,531,639	\$416,752,932
\$ Loans (total outstanding)	\$1,907,118,025	\$1,956,173,262	\$2,139,145,462	\$2,498,668,431	\$2,673,575,221	\$2,874,093,360	\$3,104,020,832
Allowance for Loan Loss	-\$23,000,537	-\$27,262,742	-\$31,063,033	-\$41,070,494	-\$43,846,634	-\$48,934,964	-\$55,640,650
Investments & Cash	\$873,079,278	\$1,258,854,789	\$1,325,962,220	\$1,502,193,955	\$1,547,343,644	\$1,640,066,214	\$1,736,569,096
Fixed & Other Assets	\$175,094,688	\$148,590,529	\$153,824,653	\$208,069,377	\$218,472,847	\$229,396,491	\$240,866,316
NCUA or ASI Deposit	\$25,754,206	\$26,491,787	\$28,439,826	\$32,419,076	\$34,445,268	\$36,598,097	\$38,885,478
Total Assets	\$2,958,045,660	\$3,362,847,625	\$3,616,309,128	\$4,200,280,345	\$4,429,990,346	\$4,731,219,198	\$5,064,701,072
Deposits							
Checking	\$639,293,995	\$619,173,024	\$654,978,280	\$776,965,415	\$825,525,753	\$879,184,927	\$938,529,910
Savings	\$1,003,250,409	\$835,456,692	\$842,616,968	\$990,922,346	\$1,047,900,381	\$1,110,774,404	\$1,180,197,804
Money Markets	\$679,632,181	\$539,182,100	\$524,960,151	\$650,447,286	\$686,221,887	\$725,679,646	\$769,220,425
Certificates	\$234,189,652	\$701,925,603	\$889,239,248	\$988,386,810	\$1,025,451,315	\$1,069,032,996	\$1,119,812,063
Other Deposits	\$122,888,887	\$202,421,729	\$199,284,539	\$161,691,422	\$161,735,608	\$168,238,893	\$182,325,316
Total Deposits	\$2,679,255,124	\$2,898,159,148	\$3,111,079,186	\$3,568,413,279	\$3,746,834,944	\$3,952,910,866	\$4,190,085,518
Other Liabilities	\$91,501,931	\$225,279,352	\$229,672,674	\$245,627,303	\$260,364,941	\$312,151,795	\$359,644,417
Net Worth	\$299,906,430	\$327,322,680	\$351,680,473	\$428,888,375	\$465,439,073	\$508,805,149	\$557,619,749
Other Equity	\$3,019,954	\$3,019,954	\$3,072,037	\$6,806,645	\$6,806,645	\$6,806,645	\$6,806,645
Unrealized Gain/Loss	-\$115,637,779	-\$90,933,509	-\$79,195,242	-\$49,455,257	-\$49,455,257	-\$49,455,257	-\$49,455,257
Total Liabilities and Equity	\$2,958,045,660	\$3,362,847,625	\$3,616,309,128	\$4,200,280,345	\$4,429,990,346	\$4,731,219,198	\$5,064,701,072

Income Statement

Income Statement - Rogue Credit Union

Formula	Dec-2022	Dec-2023	Dec-2024	Dec-2025	Dec-2026	Dec-2027	Dec-2028
Net Loan Interest Income	\$89,980,789	\$112,599,690	\$132,948,587	\$150,799,402	\$161,355,360	\$172,207,012	\$184,583,573
Investment Income	\$23,834,872	\$41,919,332	\$57,710,295	\$64,459,688	\$65,980,494	\$69,966,389	\$73,042,441
Fee & Other Operating Income	\$36,343,884	\$40,171,005	\$35,950,624	\$41,921,843	\$43,808,326	\$46,217,784	\$48,990,851
Total Gross Income	\$150,159,545	\$194,690,027	\$226,609,506	\$257,180,933	\$271,144,180	\$288,391,185	\$306,616,865
Compensation & Benefits	\$51,449,858	\$58,364,783	\$64,010,818	\$70,892,501	\$74,968,820	\$78,554,839	\$81,867,357
Travel & Conference	\$360,526	\$429,658	\$443,355	\$508,095	\$537,310	\$562,832	\$586,752
Office Occupancy	\$4,598,939	\$6,118,162	\$6,128,486	\$6,339,335	\$6,703,847	\$7,022,280	\$7,320,727
Office Operations	\$14,761,244	\$16,016,985	\$20,375,647	\$22,271,254	\$24,051,851	\$24,845,564	\$25,869,063
Marketing	\$2,890,211	\$3,500,218	\$4,680,014	\$5,437,890	\$6,975,468	\$7,110,428	\$7,425,934
Loan Servicing	\$4,117,002	\$4,008,584	\$3,830,266	\$5,427,809	\$5,739,908	\$6,012,554	\$6,268,088
Professional & Outside Services	\$8,143,767	\$11,861,659	\$11,967,622	\$14,753,855	\$17,002,202	\$16,718,307	\$17,387,898
Member Insurance & Other Operating Expenses	\$3,936,538	\$3,906,795	\$2,113,061	\$4,964,724	\$5,250,196	\$5,499,580	\$5,733,312
Total Operating Expense	\$90,258,085	\$104,206,844	\$113,549,269	\$130,595,463	\$141,229,602	\$146,326,384	\$152,459,131
Non-Op Gain/Loss	-\$9,159,621	-\$4,602,266	-\$1,131,576	\$1,849,636	\$0	\$0	\$0
Dividends	\$5,592,639	\$32,514,985	\$55,343,550	\$56,718,145	\$59,736,099	\$62,625,522	\$66,876,415
Interest on Borrowed Money	\$1,543,792	\$7,511,310	\$9,210,120	\$8,212,078	\$8,612,078	\$8,794,199	\$9,044,199
Total Cost of Funds	\$7,136,431	\$40,026,295	\$64,553,670	\$64,930,223	\$68,348,177	\$71,419,721	\$75,920,614
Provision for Loan Loss	\$8,390,021	\$17,963,033	\$23,017,198	\$23,831,596	\$25,015,703	\$27,279,004	\$29,422,520
Net Income	\$35,215,387	\$27,891,589	\$24,357,793	\$39,673,287	\$36,550,698	\$43,366,076	\$48,814,600

Key Data & Ratios

Key Data/Ratios - Rogue Credit Union

Formula	Dec-2022	Dec-2023	Dec-2024	Dec-2025	Dec-2026	Dec-2027	Dec-2028
Net Worth Ratio	10.14%	9.74%	9.72%	10.21%	10.51%	10.75%	11.01%
Return on Assets	1.19%	0.88%	0.70%	1.02%	0.85%	0.95%	1.00%
Asset Growth	0.44%	13.68%	7.54%	16.15%	5.47%	6.80%	7.05%
Deposit Growth	1.59%	8.17%	7.35%	14.70%	5.00%	5.50%	6.00%
Cost of Funds	0.27%	1.38%	2.03%	1.66%	1.58%	1.56%	1.55%
Loan Growth	13.03%	2.57%	9.35%	16.81%	7.00%	7.50%	8.00%
Loan Yield	4.94%	5.77%	6.48%	6.45%	6.24%	6.21%	6.18%
Delinquency (total)	0.42%	0.62%	0.95%	0.86%	0.90%	0.92%	0.88%
Loan Charge Offs	0.31%	0.73%	0.94%	0.74%	0.79%	0.80%	0.76%
Member Growth	9.52%	2.19%	2.87%	12.62%	4.00%	5.50%	6.00%
Loans to Shares	71.18%	67.50%	68.76%	70.02%	71.36%	72.71%	74.08%


B. 7-YEAR GRAPHS



Create bar or line charts that graphically represent trends across all 7 years (2022-2028) of data from the Balance Sheet, Income Statement and Key Ratios tables.



Since Peer-to-Peer is solely used to analyze historical information, that tool cannot be effectively utilized to produce these graphs. You will need to find an alternative means to produce them.



Specifically graph the following items and include them in your project:



- Total Assets
- Total Loans
- Total Deposits
- Equity
- Net Income
- Operating Expense
- Net Worth Ratio
- Return on Assets
- Loan to Deposits











C. FINANCIAL DATA ANALYSIS

i. Current Year Analysis

As measured by the above reports, provide a detailed narrative on observed changes in your 2025 performance compared to the previous three years, including reasons for those movements.

Rogue's 2025 financial performance represents a clear inflection point relative to the prior three-year period. Compared with 2022 through 2024, 2025 produced stronger reported earnings, improved capital strength, better asset quality, and a materially larger balance sheet. These improvements occurred at the same time as the late-year legal merger with Members 1st Credit Union, which materially increased year-end assets, loans, deposits, and capital. Importantly, however, the 2025 results reflect more than balance sheet expansion alone. Rogue entered the merger period from a position of improving underlying operating performance, with stronger recurring revenue trends, easing funding pressure, and improving credit metrics already visible within the core financials.

The balance sheet shows that 2025 growth exceeded the pace established over the prior three years. Total assets increased from \$2.96 billion in 2022 to \$3.36 billion in 2023 and \$3.62 billion in 2024, before rising to \$4.20 billion in 2025. This translated to asset growth of 16.15% in 2025, well above the 7.54% growth recorded in 2024. Total loans followed a similar pattern, increasing from \$1.91 billion in 2022 to \$2.14 billion in 2024 and then to \$2.50 billion in 2025, producing loan growth of 16.81% in 2025 compared with 9.35% in 2024. Total deposits also expanded materially, growing from \$2.68 billion in 2022 to \$3.11 billion in 2024 and then to \$3.57 billion in 2025, with deposit growth accelerating to 14.70% in 2025 from 7.35% in 2024. While these elevated growth rates reflect the late-year addition of Members 1st balances, the magnitude of change also suggests Rogue was already on a stronger growth trajectory during 2025. Loans and deposits followed the same pattern, which indicates that the merger amplified an existing trend rather than masking underlying softness. Rogue therefore exited 2025 at a substantially larger scale than any of the prior three years, with growth driven by both transaction-related expansion and improving core balance sheet momentum.

The merger with Members 1st Credit Union helps explain that year-end expansion in scale. Members 1st added meaningful balances in assets, loans, deposits, and capital, which is why the year-end growth rates in 2025 were not purely organic. At the same time, Rogue's performance cannot be explained by merger scale alone. The visible 2025 income statement shows broad-based revenue growth rather than improvement concentrated in only one category, which indicates the credit union entered the merger from a position of underlying operating strength.

Revenue performance improved materially in 2025 following compression during the prior two years. Net loan interest income increased from \$90.0 million in 2022 to \$112.6 million in 2023 and \$132.9 million in 2024, reaching \$150.8 million in 2025. This increase was driven primarily by materially higher loan balances rather than yield expansion. Although loan yield declined modestly from 6.48% in 2024 to 6.45% in 2025, that pressure was offset by improved funding conditions. Cost of funds declined meaningfully from 2.03% to 1.66%, reflecting deposit repricing dynamics and improved deposit mix. This broader spread improvement allowed Rogue to capture the earnings benefit of balance sheet growth even as asset yields stabilized, contributing to stronger net interest income and overall profitability.

Investment income also increased sharply, rising from \$23.8 million in 2022 to \$41.9 million in 2023, \$57.7 million in 2024, and \$64.5 million in 2025. This increase is consistent with elevated liquidity levels and favorable yields within the investment portfolio. Fee and other operating income declined in 2024 but recovered in 2025 to \$41.9 million. Together, these categories drove total gross income to \$257.2 million in 2025, up from \$226.6 million in 2024 and \$150.2 million in 2022. The improvement in unrealized losses during the year further supports the quality of this performance, indicating reduced market pressure on capital relative to 2024. In addition, 2025 included a positive non-operating gain of \$1.85 million, which strengthened reported earnings for the year.

Strategic product activity in 2025 also reinforces the financial narrative. Rogue introduced *MyRewards Checking* and *MyMoney Market* as participation-focused deposit products designed to deepen primary financial institution relationships, improve deposit stability, and support lower funding costs over time. Later in the year, we launched the *Visa Signature Rewards* credit card to strengthen long-term revenue potential through interchange income while deepening relationships with higher-quality borrowers. Although these initiatives had varying degrees of immediate impact in 2025, they reflect intentional investments aligned with future margin stability, revenue diversification, and relationship-driven growth.

Expense growth remained significant in 2025, but its composition suggests expansion tied to scale, product development, and operational complexity rather than deterioration in performance. Total operating expense rose from \$113.5 million in 2024 to \$130.6 million in 2025. Compensation and benefits increased from \$64.0 million to \$70.9 million, reflecting a larger and more complex organization. Loan servicing expense increased from \$3.8 million to \$5.4 million, professional and outside services rose from \$12.0 million to \$14.8 million, and member insurance and other operating expenses increased from \$2.1 million to \$5.0 million. Marketing expense also rose from \$4.7 million to \$5.4 million, consistent with product launches and broader strategic activity. Even with that higher cost base, revenue growth outpaced expense growth during the year, which indicates Rogue realized positive operating leverage in 2025. Rogue still entered 2026 with a higher expense base that will require continued discipline, but the 2025 pattern does not suggest weakening core performance.

Credit provisioning increased alongside growth, while asset quality improved. Provision for loan loss increased from \$23.0 million in 2024 to \$23.8 million in 2025, which is consistent with continued loan expansion and prudent reserving. At the same time, total delinquency improved from 0.95% in 2024 to 0.86% in 2025, and net charge-offs improved from 0.94% to 0.74%. Rogue, therefore, entered the merger period from a position of stronger credit quality rather than emerging weakness.

Overall, reported profitability rebounded strongly in 2025. Net income increased from \$24.4 million in 2024 to \$39.7 million in 2025, reversing the decline experienced from 2022 through 2024. Return on assets improved from 0.70% to 1.02%. Capital strength also improved materially. Net worth increased from \$351.7 million in 2024 to \$428.9 million in 2025, and the net worth ratio rose from 9.72% to 10.21%. While reported net income benefited from the one-time non-operating gain, core operating performance also improved meaningfully through higher net interest income, lower funding costs, stable credit provisioning, and improving asset quality. As a result, the improvement in earnings cannot be attributed solely to non-recurring items.

Overall, Rogue Credit Union finished 2025 in a substantially stronger financial position than it held during 2022 through 2024. The clearest reading of 2025 is that Rogue improved reported earnings, strengthened capital, improved credit quality, and closed the year at a much larger scale. The merger materially increased year-end assets, loans, deposits, and capital, while the income statement and key ratios show that Rogue also generated stronger revenue and stronger reported profitability during the year. Rogue therefore enters 2026 with greater scale, stronger capital, and improved earnings power, alongside the challenge of sustaining performance after the one-time non-operating gain falls away.

ii. Forecast Assumptions & Analysis

Provide a detailed explanation of the rationale for your forecasted assumptions.

The process of developing and implementing goals and then forecasting their results is certainly an inexact science. Therefore, it is important to step back and take a look at the outcome of those forecasts from a broader perspective once they have been assembled.

Consider the following in your analysis: Does the resulting forecast make sense? Are the results consistent with your narratives? Are the key data and ratios in line with board and management expectations or risk tolerance levels? Do the graphs show the trends you expected? Is each one increasing/decreasing in the manner that you anticipated? Why or why not? What other actions might need to be considered to better align the forecasted results with your initial expectations?

The three-year forecast was developed to test the financial outcome of the strategic goals proposed in this project, with one central question in mind: what results could Rogue achieve if it invested meaningfully in AI-enabled automation, digital capability, and operational efficiency, not to reduce FTE, but to redeploy employee capacity toward higher-value member contact, stronger relationship development, and better service? Based on that strategy, I built the forecast to assume near-term investment pressure in exchange for longer-term gains in productivity, growth, and operating performance. I did not introduce a separate macroeconomic shock, such as a major recession, interest rate disruption, or regulatory event, as the primary driver of the forecast. Instead, the model is driven primarily by the expected impact of the five proposed strategic goals and by the operational reality of the November 2025 merger, including expansion into the Redding and broader Northern California market. Each of the five goals was intentionally tied to a specific financial lever in the forecast: Goal 1 to operating efficiency, Goal 2 to loan and deposit growth, Goal 3 to service delivery costs and engagement, Goal 4 to product and process performance, and Goal 5 to capital strength and profitability. That starting point matters because 2025 is not a perfectly clean year for comparison. The merger created a new combined base for 2026, and 2025 also included a positive non-operating gain that strengthened reported earnings. For that reason, I treated 2026 as the first full baseline year for the combined organization and then evaluated how my strategic goals should shape performance from that point forward.

The first major forecast driver is Goal 1: Enterprise AI-Enabled Operational Efficiency. This goal was designed to reduce manual processing, improve productivity, and moderate operating expense growth over time through AI-enabled automation, decision-support tools, workflow redesign, and knowledge integration. The expected financial result of this goal was not an immediate drop in expenses in 2026, but a slower rate of expense growth over time and gradual improvement in efficiency as automation matured. That is generally what the forecast shows. Total operating expense rises from

approximately \$130.6 million in 2025 to \$141.2 million in 2026, then to \$146.3 million in 2027 and \$152.5 million in 2028. That translates to growth of roughly 8.1% in 2026, 3.6% in 2027, and 4.2% in 2028. In other words, the goal works directionally because expense growth moderates after the heavy 2026 transition year, but it does not produce immediate, dramatic savings. I think that is the most realistic interpretation. The strategy appears sound, but the benefits scale more gradually than the original target implied. The cost of implementation, integration, change management, and training shows up first; the payoff follows later.

The second major forecast driver is Goal 2: Strengthen Primary Financial Institution Relationships Through Member Participation. This goal has the clearest direct connection to the balance sheet because it was intentionally designed to deepen participation across checking, lending, savings, and digital payment services. The stated financial logic behind the goal was that stronger participation would support approximately 7% to 8% annual loan growth and 5% to 6% annual deposit growth during the forecast period. The forecast reflects that directly. Loan growth is projected at 7.0% in 2026, 7.5% in 2027, and 8.0% in 2028, while deposit growth is projected at 5.0%, 5.5%, and 6.0%. Total loans grow from about \$2.50 billion in 2025 to \$2.67 billion in 2026, \$2.87 billion in 2027, and \$3.10 billion in 2028. Total deposits rise from about \$3.57 billion to \$3.75 billion, \$3.95 billion, and \$4.19 billion over the same period. This tells a coherent story: if more households use Rogue as their primary institution, loan demand strengthens, deposits become more stable, and relationship value increases over time. The loan-to-deposit ratio also rises gradually from about 70.0% in 2025 to 71.4%, 72.7%, and 74.1%, suggesting Rogue is putting more deposits to work without moving into a liquidity position that appears overly aggressive.

The third driver is Goal 3: Expand Digital and Data Capabilities to Support Personalized Member Experiences. This goal was designed to increase digital self-service, improve analytics, and support more personalized member engagement. Financially, this goal reinforces Goal 2 by helping

deepen relationships and improve product usage, but it also supports Goal 1 by gradually shifting routine transactions away from higher-cost service channels. That is why the goal matters even if it does not create one dramatic line-item change by itself. The forecast supports this rationale. Fee and other operating income increases from approximately \$41.9 million in 2025 to \$43.8 million in 2026, \$46.2 million in 2027, and \$49.0 million in 2028. ROA also improves gradually from 0.85% in 2026 to 0.95% in 2027 and 1.00% in 2028, suggesting that stronger digital adoption and better member engagement contribute to improved profitability over time. This result is consistent with the goal's stated financial rationale that digital capability would help support stronger participation, better service delivery, and more efficient operating performance.

The fourth driver is Goal 4: Strengthen Enterprise Process Governance and Deposit Product Ownership. This goal is less flashy than the AI and participation goals, but it plays an important supporting role in the model. Its purpose was to improve operational consistency, product accountability, pricing discipline, and scalability as Rogue continues to grow. Financially, I expected this goal to support deposit product performance, more effective decision-making, and better funding discipline. The forecast shows that effect most clearly through healthy deposit growth and declining funding pressure. Deposit growth remains positive across all three forecast years, while cost of funds improves from 1.66% in 2025 to 1.58% in 2026, 1.56% in 2027, and 1.55% in 2028. Checking, savings, and money market balances all continue to rise over the period, which suggests the forecast is not relying solely on higher-cost funding to support growth. Instead, it reflects stronger core deposit performance and better product discipline, which is exactly what this goal was intended to support. This goal does not produce one headline ratio by itself, but it strengthens the infrastructure behind several of them. Sometimes strategy gets the spotlight, and sometimes strategy fixes the plumbing. Both matter. The fifth driver is Goal 5: Maintain Financial Strength While Funding Strategic Innovation. This goal served as the guardrail for the entire forecast. Rather than maximizing short-term earnings, it required

the model to maintain strong capital and profitability while still funding technology modernization and member experience improvements. This goal appears to work clearly in the forecast. Net income declines from the unusually strong 2025 level of approximately \$39.7 million to \$36.6 million in 2026, then improves to \$43.4 million in 2027 and \$48.8 million in 2028. ROA follows the same pattern, moving from 1.02% in 2025 to 0.85% in 2026 before recovering to 0.95% in 2027 and 1.00% in 2028. This pattern is consistent with the strategy: 2026 functions as a transition and investment year rather than an immediate payoff year. The net worth ratio also remains strong throughout the forecast, increasing from 10.21% in 2025 to 10.51% in 2026, 10.75% in 2027, and 11.01% in 2028. That confirms the forecast preserves capital strength while still allowing for growth and strategic investment. At the same time, ending slightly above my preferred upper range suggests the forecast still leans somewhat conservative. Apparently, when left unsupervised, I gave Rogue a slightly thicker financial life jacket than planned. While that supports safety and soundness, it may also suggest the forecast does not deploy capital quite as aggressively as originally envisioned.

Looking at the forecast as a whole, the year-by-year pattern makes sense and is consistent with the strategic narrative. 2026 functions as a transition and investment year. Growth remains solid, but expense pressure is still visible because the organization is absorbing implementation costs, elevated marketing related to merger onboarding and cross-sell, and the normal friction that comes with integration and expansion. Marketing expense, for example, rises from approximately \$5.4 million in 2025 to \$7.0 million in 2026, which is consistent with the assumption that onboarding, retention, and relationship-deepening efforts would be elevated in the first full year after the merger. 2027 represents the first stronger realization year, with faster membership, loan, deposit, and fee growth, along with moderating expense growth and improved ROA. 2028 shows the fuller payoff year, with stronger profitability, continued capital strength, and better overall operating performance. Asset growth also follows the expected pattern, moving from 16.15% in 2025 to 5.47% in 2026, 6.80% in 2027, and 7.05%

in 2028, with total assets increasing from about \$4.20 billion in 2025 to \$4.43 billion, \$4.73 billion, and \$5.06 billion. Overall, the major graphs should show the trends I expected: upward movement in loans, deposits, assets, fee income, net income, and net worth, with profitability and capital strengthening after the initial investment year.

From a broader risk perspective, the forecast remains reasonable and generally in line with board and management expectations for safety and soundness. ROA returns to 1.00% by 2028, the net worth ratio remains comfortably above 10%, and the loan-to-deposit ratio increases gradually without signaling undue liquidity strain. Loan yield declines modestly from 6.24% in 2026 to 6.18% in 2028, which is believable in a competitive lending environment where Rogue prioritizes relationship growth and retention over aggressive repricing. Delinquency and charge-offs remain elevated but manageable, which is also believable in a forecast that includes post-merger integration and continued growth. The one area where the model does not fully achieve my original expectation is efficiency. Because expense pressure moderates more slowly than initially hoped, the forecast suggests that benefits realization from AI, automation, and digital enablement may take longer than the original goal implied.

If actual performance were to lag the forecast, the right response would depend on which strategic lever was underperforming. If loan and deposit growth come in below expectations, Goal 2 assumptions around participation may need to be revisited. If operating expenses do not moderate as projected, Goals 1 and 3 may take longer to translate into measurable savings. If deposit mix or product performance remains weaker than expected, Goal 4 may require stronger implementation discipline. If the net worth ratio remains above the preferred range for too long, management may need to consider whether capital is being deployed aggressively enough through faster asset growth, additional strategic investment, or greater short-term earnings pressure in support of long-term value.

Overall, the forecast is credible because it reflects the expected financial effect of my proposed goals rather than simply describing what the spreadsheets happen to show. Goal 2 has the strongest

direct impact on growth, while Goals 1 and 3 improve performance more gradually by moderating expense pressure and supporting better profitability over time. Goal 4 strengthens the organizational discipline behind product and deposit performance, and Goal 5 ensures all of those efforts remain within reasonable capital and profitability boundaries. As a result, the forecast reflects a deliberate strategic tradeoff: Rogue accepts heavier expense pressure in 2026 in order to invest in automation, digital capability, onboarding, and process redesign that support stronger growth and better operating performance over time. This cause-and-effect relationship is what makes the forecast believable. It shows a realistic sequence of investment first, stronger realization second, and fuller payoff third, while also acknowledging that not every strategic goal produces a perfect financial result on the same timeline.

VII. SPREAD ANALYSIS

Using your Balance Sheet and Income Statement for 2025 as well as your forecast for 2028, you will now complete the Spread Analysis table below. This exercise will provide insights into your credit union's performance in the five different areas or "levers" that determine profitability. You will display the completed table below in your project. Once the table is complete, you are expected to write an analysis of the results.

The dollar amounts can be taken directly from your Income Statements as follows:

- Interest Income: Loan Interest Income + Investment Interest Income
- Cost of Funds: Total Cost of Funds
- Net Interest Margin: Interest Income – Cost of Funds
- Operating Expense: Total Operating Expense
- Provision for Loan Loss: Provision for Loan Loss
- Fee & Other Income: Fee & Other Income + Non-Operating Gains/(Losses)
 - [Note: Add the non-operating amount if it is a gain; subtract it if it is a loss]
- Net Income (ROA): Net Interest Margin – Operating Expense – Provision for Loan Loss + Fee & Other Income

The percentages are calculated by dividing the dollar amount for each line item by Average Assets. "Average Assets" is calculated as follows:

- 2025 Average Assets: (2024 Total Assets + 2025 Total Assets) / 2
- 2028 Average Assets: (2027 Total Assets + 2028 Total Assets) / 2

Spread Analysis

		2025\$	2025%	2028\$	2028%
	Interest Income	\$215,259,090	5.51%	\$257,626,014	5.26%
-	Cost of Funds	\$64,930,223	1.66%	\$75,920,614	1.55%
=	Net Interest Margin	\$150,328,867	3.85%	\$181,705,400	3.71%
-	Operating Expense	\$130,595,463	3.34%	\$152,459,131	3.11%
-	Provision for Loan Loss	\$23,831,596	0.61%	\$29,422,520	0.60%
+	Fee & Other Income	\$43,771,479	1.12%	\$48,990,851	1.00%
=	Net Income (ROA) *	\$39,673,287	1.02%	\$48,814,600	1.00%

Analyze Your Results

Consider the following questions: What changes in profitability (ROA) does your forecast project? Which of the 5 components of Net Income are driving these changes and why? Are these components increasing/decreasing the way you anticipated? What business factors are impacting these changes?

For purposes of this analysis, 2025 average assets were calculated using the average of 2024 and 2025 total assets, while 2028 average assets were calculated using the average of 2027 and 2028 total assets. Using that approach, Rogue's return on assets moves from 1.02% in 2025 to 1.00% in 2028. That represents only a slight decline in reported profitability on a spread basis and still reflects strong earnings performance by the end of the forecast period. This result is reasonable because 2025 includes a one-time positive non-operating gain and reflects an unusually strong post-merger year, while 2028 reflects a more normalized earnings environment.

The strongest positive driver of profitability in the forecast is the operating expense ratio. Total operating expense increases in dollars from \$130.6 million in 2025 to \$152.5 million in 2028, but declines from 3.34% to 3.11% of average assets. This indicates Rogue is gaining operating leverage as it grows. That pattern is consistent with the strategic direction of the forecast, which assumes the credit union will improve efficiency through process redesign, stronger workflows, digital tools, and AI-enabled support capabilities. The benefit is not based on reducing service levels or broadly cutting staff. It comes from scaling the organization more effectively as assets, loans, and deposits continue to grow.

Interest income also increases meaningfully in dollars, rising from \$215.3 million in 2025 to \$257.6 million in 2028. This reflects continued balance sheet growth and stronger loan production over the forecast period. Even so, interest income declines from 5.51% to 5.26% of average assets, which suggests the forecast is not relying on a more favorable yield environment to drive earnings. Instead, the model assumes Rogue will generate more income primarily through growth in loans and investments rather than through materially richer pricing. That result fits a forecast built around participation growth, relationship deepening, and post-merger expansion rather than dependence on easier external conditions.

Cost of funds improves modestly from 1.66% in 2025 to 1.55% in 2028. This supports earnings and suggests stronger funding discipline as the forecast progresses. Even with that improvement, the decline in interest income as a percentage of average assets is enough to create some spread pressure. As

a result, net interest margin declines from 3.85% to 3.71%. Margin compression is one of the main reasons profitability softens slightly by 2028. It also makes the forecast more believable. Rogue is not assumed to operate in a magically easier environment. Instead, the spread analysis shows that it must preserve earnings through disciplined growth, better operating leverage, and continued control of funding costs.

Provision for loan loss remains relatively stable, moving from 0.61% in 2025 to 0.60% in 2028. This indicates that credit costs remain controlled across the forecast period and do not serve as a major source of either earnings improvement or earnings deterioration. Given Rogue's continued growth, post-merger integration, and expansion assumptions, this is a reasonable outcome. The forecast does not assume unusually low credit costs, but it also does not project deterioration significant enough to materially disrupt profitability.

Fee and other income increases in dollars from \$43.8 million in 2025 to \$49.0 million in 2028, but declines from 1.12% to 1.00% of average assets. That makes fee and other income another mild headwind to profitability on a spread basis. The dollars improve, but the ratio declines because asset growth outpaces growth in this revenue source. This suggests the forecast is not dependent on fee expansion to sustain earnings. That result is consistent with a member-focused strategy in which financial improvement is expected to come more from deeper relationships, stronger participation, and operating efficiency than from increasing reliance on non-interest income.

The spread analysis shows Rogue's forecasted profitability being shaped by two opposing forces. The clearest positive influence is improved operating efficiency, supported by growth in interest income dollars, modest improvement in cost of funds, and stable credit costs. The clearest negative influences are margin compression and the decline in fee and other income as a percentage of average assets. Because those pressures offset part of the benefit from efficiency improvement, the forecast projects ROA at 1.00% in 2028 rather than a stronger increase from the 2025 level. Even so, the decline is

minimal, and the broader message is that Rogue appears able to preserve strong profitability while supporting continued growth, integration, and long-term financial strength.

The results move largely in the direction I expected, although not perfectly in every category. I expected operating efficiency to become one of the clearest financial benefits of the strategy, and that is evident in the declining operating expense ratio. I also expected continued loan and deposit growth to support higher income in dollars, and that is reflected in the growth of interest income. The areas that did not improve as much as I initially hoped were net interest margin and fee and other income as a percentage of average assets. Those two factors create enough pressure to keep ROA roughly flat by 2028. Even so, the forecast still supports the conclusion that Rogue's strategic initiatives should help preserve strong profitability over time while supporting continued growth, post-merger integration, and long-term financial strength.

VIII. FEASIBILITY ASSESSMENT ON IMPLEMENTATION

Now that you have completed your analysis, proposed five strategic goals, and developed a financial forecast, this final step asks you to assess the realistic potential for implementation. This is your opportunity to evaluate how well your plan aligns with your credit union's current direction, capacity, and culture—and to reflect on what it would take to move your ideas from concept to execution.

Your response should offer both practical insight and personal perspective. Consider the operational, political, and cultural realities within your organization, and how they may influence whether your proposed strategies are embraced.

Use the questions below to guide your assessment:

- What will it take for your credit union to actually adopt all or part of your 5-goal plan?
- In your opinion, which components would have the greatest chance of being adopted by your credit union?
- In your opinion, which components would have the least chance of being adopted by your credit union?
- What factors lead you to have that opinion?
- How do your goals differ from your credit union's current goals? How are they similar?

This section should reflect thoughtful analysis and authentic reflection. There are no right or wrong answers, only your reasoned, experience-based assessment of feasibility and fit.

Based on my understanding of Rogue's current direction, culture, and capacity, I believe the goals proposed in this project are directionally aligned with where Rogue is already headed. They are not disconnected or theoretical ideas. They fit many of the priorities Rogue already values. At the same time, feasibility is not just a question of whether a goal makes strategic sense. It is also a question of how that goal would move through leadership discussion, prioritization, funding, and cross-functional ownership. That is where implementation becomes more complicated. At Rogue, as in most organizations, the challenge is often less about whether the ideas are strong and more about time, focus, ownership, and execution capacity.

For Rogue to adopt all or part of this plan, it would require clear business cases, strong executive sponsorship, defined ownership, and realistic sequencing. The goals would need to be translated into measurable outcomes, implementation steps, and practical timelines. Several of these goals, especially those related to AI-enabled efficiency, digital capability, and enterprise process governance, cross

multiple departments and would be difficult to move forward as isolated initiatives. In my opinion, the central challenge is not whether Rogue could do these things. It is whether Rogue would be able to prioritize them appropriately, assign clear ownership, and create enough space to execute them well alongside other work already in progress.

The most realistic path to adoption would likely be through Rogue's existing leadership and prioritization channels rather than as a separate package of five goals. In practice, that means the goals with the clearest strategic alignment, visible operational value, and more obvious ownership would likely move more easily, while the goals that require broader cross-functional governance or more significant organizational change would likely move more slowly. That distinction matters because it reflects how implementation usually works in real organizations. Not every strong idea moves at the same speed, and not every idea requires the same level of sponsorship or coordination.

In my opinion, the components with the greatest chance of adoption are the ones that most clearly extend work Rogue is already moving toward. These include strengthening primary financial institution relationships through participation, expanding digital engagement and self-service capabilities, and maintaining financial strength while funding innovation. These goals feel like a natural fit because they align with priorities Rogue already values and because they produce visible, measurable outcomes. I also believe parts of the AI-enabled operational efficiency goal have a good chance of adoption, especially where they build on work Rogue is already doing. The strongest aspect of that goal is that it frames automation as a way to support employees and improve service rather than replace people. That matters culturally. Rogue is relationship-based, and I do not believe a headcount-reduction narrative would fit nearly as well as a capacity-redeployment narrative.

I also believe Rogue would likely be receptive to enterprise process governance and deposit product ownership in concept, but slower to fully formalize it. I see this as a strong-fit goal because it addresses real operational complexity and would help the organization scale more effectively. As Rogue grows,

clearer ownership and better process discipline become more important. Where I think implementation becomes more difficult is when it is fully formalized. Ownership models can shift decision-making, accountability, and existing ways of working, which makes implementation more complex even when the concept itself is well aligned. This is not a poor-fit goal. It is a good-fit goal that may move more slowly once real ownership questions surface.

In my opinion, the components with the least chance of immediate adoption are the broadest, most enterprise-wide parts of the plan if they are presented as one large transformation rather than as more targeted work. That is especially true for a full-scale AI and personalization strategy. I do believe Rogue, in general, will need to continue strengthening digital capability, analytics, and automation. That is simply the direction of the business and the broader financial services environment. What I am less certain about is whether Rogue would choose to move all of those pieces forward at the same pace without meaningful piloting, sequencing, and proof along the way. In practice, enterprise-wide ideas usually require stronger sponsorship, more cross-functional alignment, and clearer governance before they are ready to scale. Change is usually more straightforward on paper than it is in practice, especially when ownership crosses multiple teams.

Several factors led me to that conclusion. Rogue's culture values service quality, collaboration, and employee engagement, so initiatives that improve the member experience and help employees do their jobs better are more likely to gain traction than initiatives framed primarily around efficiency alone. Rogue is also a regulated financial institution, which means there is an appropriate level of caution around emerging technology, data use, and automation. Change capacity is real as well. The more cross-functional, technical, and politically sensitive a goal becomes, the more sponsorship, coordination, and sequencing it will require. In practice, that means some goals would likely move through leadership more quickly because their ownership and value are easier to define, while others would require more

discussion to determine who leads the work, how it is governed, and where it fits among competing priorities.

My goals are similar to Rogue's current goals in several important ways, but they also differ somewhat in framing and emphasis. They reinforce themes Rogue already values, including participation, digital engagement, operational improvement, and financial strength. In many ways, they extend work Rogue is already doing or already discussing. The biggest difference is that, in this project, I made the connection between strategy and measurable financial outcomes more explicit in order to support the financial forecast and analysis. I also placed greater emphasis on formal ownership, process discipline, and the use of automation and AI as enterprise levers rather than simply supportive tools operating in the background. In that sense, the goals are not inconsistent with Rogue's direction, but they are more specifically framed for the purposes of this project.

If I had to summarize feasibility in one sentence, I would say this: these goals fit Rogue directionally, but successful implementation would depend on how effectively they move through leadership sponsorship, prioritization, and cross-functional ownership. Based on my experience, I believe Rogue could move in this direction, especially where the goals clearly support member service, operational discipline, and long-term growth. At the same time, I also know implementation depends on more than a good idea. It depends on timing, ownership, and making enough room to do the work well. I believe the goals are feasible if they are introduced in a way that respects Rogue's culture, capacity, and appetite for change.

IX. EXECUTIVE SUMMARY

Write a clear and concise Executive Summary that synthesizes the most important findings, decisions, and recommendations from your entire project. This summary should serve as a standalone overview—something a senior leader could read to quickly understand the key takeaways and strategic direction you’ve proposed, without needing to read the full report.

- The Executive Summary should be 1–2 pages in length and written after the rest of your project is complete.
- Think of this as your opportunity to clearly communicate the value of your work to someone who hasn’t seen the details.
- Your Executive Summary should include the following components:
- A brief description of your credit union’s current position, drawing from your financial and operational analysis
- A summary of the major external trends and inflection points likely to impact the credit union
- Key insights from your SWOT analysis (especially those that shaped your planning)
- An overview of your five strategic goals, including what they aim to achieve and how they align with broader strategy
- Highlights from your three-year financial forecast, including expected outcomes and major investments or risks
- A short reflection on the feasibility and strategic fit of your proposed plan

Use clear, professional language. Avoid repeating the full content of your project—focus instead on synthesizing and communicating the most critical insights and strategic decisions.

Your Executive Summary should reflect the tone and format you would use in presenting this work to your board or senior leadership team.

Rogue Credit Union enters the 2026–2028 planning period from a position of strength, but also at a meaningful point of transition. During 2025, Rogue improved reported profitability, strengthened capital, improved asset quality, and ended the year at a materially larger scale following the legal merger with Members 1st Credit Union. Total assets increased to approximately \$4.20 billion, net income rose to approximately \$39.7 million, return on assets improved to 1.02%, and the net worth ratio increased to 10.21%. Asset quality also improved, with total delinquency declining to 0.86% and net charge-offs improving to 0.74%. At the same time, the merger expanded Rogue’s geographic reach and introduced additional operational complexity, making 2026 the first full year in which Rogue must balance integration demands with continued growth and strategic execution.

This project finds that Rogue’s strategic environment is being shaped by five major external inflection points: artificial intelligence, evolving consumer expectations, fragmentation of financial relationships, demographic change, and increased competition from non-traditional financial providers. Together, these forces are changing how members access financial services and how institutions compete for relevance. Members increasingly expect seamless digital self-service, personalized support, and convenient access across channels, while fintechs and embedded finance providers make it easier for consumers to spread their financial lives across multiple organizations. In this environment, Rogue’s long-term success will depend not only on technology adoption, but on its ability to use technology in ways that strengthen rather than weaken its relationship-based service model. Artificial intelligence is particularly important because it creates an opportunity to improve operational efficiency while enabling employees to provide more timely, informed, and personalized support.

The SWOT analysis suggests that Rogue is well-positioned to respond to these changes, but not without discipline. Its most important strengths include a strong relationship-based service culture, clear enterprise strategic alignment, growing knowledge-management and digital-enablement capabilities, a participation-based relationship model, and solid financial capacity. At the same time, Rogue faces

several internal challenges, including uneven enterprise process maturity, increasing operational complexity from growth and merger activity, the challenge of absorbing technology consistently across the enterprise, and the need for more formal deposit product ownership and governance. Externally, Rogue has meaningful opportunities to deepen household relationships, expand digital and data-enabled service delivery, and use AI to improve employee productivity and member experience. However, those opportunities exist alongside real threats, including rising digital expectations, stronger fintech competition, deposit pressure, and the risk of cultural dilution during continued expansion.

Based on these findings, this project proposes five strategic goals. The first is to implement AI-enabled automation and decision-support tools across key operational and service functions in order to reduce manual processing, improve productivity, and moderate expense growth over time. The second is to deepen Rogue's role as members' primary financial institution by increasing participation across checking, lending, savings, and digital payment services. The third is to expand digital and data capabilities to support more personalized member experiences and greater self-service adoption. The fourth is to strengthen enterprise process governance and formalize deposit product ownership in order to improve consistency, accountability, and product performance. The fifth is to maintain financial strength while funding innovation, ensuring Rogue remains well-capitalized and profitable even as it invests in technology modernization and member experience improvements.

The three-year financial forecast indicates that these goals are directionally sound and financially feasible. The model reflects 2026 as a transition and investment year, with ongoing merger onboarding, elevated operating expense, and continued strategic spending tied to automation, digital capability, and process improvement. By 2027 and 2028, the financial benefits become more visible as growth strengthens and expense pressure moderates. Total assets are projected to increase from approximately \$4.20 billion in 2025 to \$5.06 billion by 2028. Loans are projected to grow to approximately \$3.10 billion, while deposits are projected to increase to approximately \$4.19 billion. Net income is projected

to move from approximately \$36.6 million in 2026 to \$48.8 million in 2028, while return on assets improves from 0.85% to 1.00%. The net worth ratio remains strong throughout the forecast period, rising from 10.51% in 2026 to 11.01% in 2028. This supports safety and soundness, while also suggesting that the forecast remains somewhat conservative in its deployment of capital.

Overall, this project concludes that Rogue is strategically well-positioned for continued growth, provided it can execute with discipline. The proposed goals fit Rogue's direction, culture, and financial capacity, but enterprise-level initiatives will require strong sponsorship, clear ownership, and thoughtful sequencing to succeed. The plan is most compelling not because it assumes technology will replace people, but because it assumes Rogue can use technology, especially AI and digital enablement, to reduce routine work and redirect employee capacity toward higher-value member contact, stronger relationship development, and better service. In that sense, the strategy is both financially credible and culturally aligned. It supports growth, protects capital, and reinforces Rogue's long-term differentiation as a credit union that combines operational strength with human-centered member service.

X. YOUR CONCLUSIONS

This final section is an opportunity for you to reflect on your personal growth, professional development, and overall experience completing the project. Your responses should be thoughtful, specific, and authentic. You are encouraged to draw connections between the skills you applied, the insights you gained, and how this experience may influence your work moving forward. You can write this as one complete narrative or respond to the prompts one at a time.

Please respond to each of the following prompts:

Personal and Professional Impact

- How has completing this project influenced the way you approach your role, your decision-making, or how you set and pursue personal and departmental objectives?

Project Management Reflection

- What did you learn about your own project management skills throughout this process? Consider aspects such as time management, planning, organization, adaptability, or collaboration.

Most Enjoyable and Most Challenging Components

- Which section of the project did you most enjoy and why? Which section was most challenging for you, and what specific steps will you take to build your skills or confidence in this area?

Recommendations for Improvement

- Do you have any suggestions for improving this project experience for future participants? Are there topics or skills you wish had been covered in the curriculum that would have helped you complete the project more effectively?

Personal and Professional Impact

Completing this project had a bigger impact on me than I expected. On one hand, it reminded me that I actually do have a strong writing voice. I spent years working with English professors, and this project gave me a chance to lean into that part of myself again. It made me take pride in my writing style and in my ability to organize ideas, connect themes, and tell a larger strategic story. On the other hand, there were definitely moments when this project made me question all of my life choices and wonder whether I should just go work the checkout line at Fred Meyer. Honestly, that is still kind of my dream job, so it was not a terrible backup plan.

More seriously, this project changed the way I think about my role and my decision-making. It pushed me to think beyond my own department and view decisions through a broader enterprise lens. I had to consider not only what made sense operationally, but also what aligned with the credit union's long-term strategy, member experience, and financial realities. It reinforced that good ideas are not enough on their own. They have to be connected to larger organizational goals and supported by realistic assumptions. That is something I will carry forward in how I set personal goals, support departmental objectives, and contribute to larger conversations at Rogue.

Project Management Reflection

This project taught me a lot about my own project management skills. I learned that one of my strengths is being able to pull together a large amount of information and shape it into something more organized and meaningful. I also learned that I do my best work when I can connect ideas across sections rather than treating each part like a separate assignment.

One of the best parts of the process was the opportunity to talk with my fellow Psi classmates and with others here at Rogue about the project. Those conversations were encouraging, helpful, and honestly just made the process feel less lonely. It was reassuring to hear how others were approaching the same challenges, and it reminded me that sometimes project management is not just about timelines and organization. Sometimes it is about collaboration, talking things through, and realizing you are not the only person staring at a spreadsheet like it personally offended you.

This project also showed me that I can improve in pacing and confidence. Some sections took much more revisiting than I expected, especially as my thinking became more refined. I learned that adaptability matters just as much as planning. Sometimes the work is not linear, and part of managing a project well is allowing room to rethink, revise, and make the final product stronger.

Most Enjoyable and Most Challenging Components

The section I most enjoyed was the strategic goals portion of the project. I really like building goals, and working on that section reminded me of one of my favorite jobs at Rogue Community College, where I helped develop college goals and assess them for accreditation. There is something very satisfying to me about taking big ideas and turning them into clear, purposeful objectives. It feels like the place where vision becomes action. That section felt especially natural to me because it required both strategic thinking and structure, which is a combination I genuinely enjoy.

The most challenging component for me was the financial forecasting section. That part stretched me the most because it required me to translate strategic ideas into numbers and assumptions that had to make sense together. It was a completely different kind of thinking from writing narrative sections, and at times it felt intimidating. There were definitely moments when I felt confident and other moments when I felt like the spreadsheet was winning. Even so, it was one of the most valuable parts of the project because it forced me to think more deeply about how strategy, growth, risk, and financial performance all connect.

To continue building my confidence in this area, I plan to spend more time strengthening my financial forecasting and analysis skills. I also want to keep learning from finance and accounting partners and get more comfortable not only working through the numbers, but also explaining the assumptions behind them in a clear and credible way.

Recommendations for Improvement

One suggestion I would offer for improving the project experience is to spend more time covering financial forecasting in class. For me, that was the area where additional instruction would have been the most helpful. The financial section is such an important part of the project, but it can also be one of the hardest sections to translate from concept to execution. More classroom discussion,

examples, or guided practice around building assumptions and connecting the strategy to the forecast would make the experience even more valuable for future participants.

Overall, this project was a meaningful growth experience for me. It challenged me to think more strategically, write more intentionally, and stretch into areas where I did not feel naturally confident. It also reminded me that growth usually looks a lot less polished in the middle than it does at the end.

XI. REFERENCES

- Accenture. (2023). *Global banking consumer study: The digital transformation of financial services*. <https://www.accenture.com/content/dam/accenture/final/industry/banking/document/Accenture-Banking-Consumer-Study.pdf>
- Achieva Credit Union. (n.d.). *Our story*. <https://www.achievacu.com/Home/OurStory>
- Achieva Credit Union. (2024). *Annual report*. <https://www.achievacu.com/pdfs/AnnualReport.pdf>
- American Bankers Association. (2023, October 26). *ABA survey: Mobile apps remain primary vehicle for accessing banking services*. <https://bankingjournal.aba.com/2023/10/aba-survey-mobile-apps-remain-primary-vehicle-for-accessing-banking-services/>
- Ascend Federal Credit Union. (2025a). *Annual report*. <https://image.storage.ascend.org/craft-production/pdfs/Annual-Report-PDFs/A2120-Annual-Report-8.5x5.5-Q1FEB25.pdf>
- Ascend Federal Credit Union. (2025b, July 21). *Practical financial tips for students*. <https://ascend.org/blog/practical-financial-tips-for-students>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (Working Paper No. 31161). National Bureau of Economic Research.
- Callahan & Associates. (2025). *Peer-to-Peer comparison tool*. Retrieved January 10, 2026, from Callahan portal.
- COMPASS. (2024). *2024 population estimates*. https://compassidaho.org/wp-content/uploads/Revised_population_estimates_2024.pdf
- Consumer Financial Protection Bureau. (2023, June 1). *CFPB finds that billions of dollars stored on popular payment apps may lack federal insurance*. <https://www.consumerfinance.gov/about-us/newsroom/cfpb-finds-billions-of-dollars-stored-on-popular-payment-apps-may-lack-federal-insurance/>
- CUInsight. (2022, September 30). *OCCU uses machine learning to reduce holds on deposits*. <https://www.cuinsight.com/press-release/occu-uses-machine-learning-to-reduce-holds-on-deposits>
- Curinos. (2023). *Achieva Credit Union selects Curinos to optimize deposit strategy*. <https://curinos.com/press/press-releases/achieva-credit-union-selects-curinos-to-optimize-members-saving-experience-amidst-changing-rate-environments>
- Deloitte. (2024). *2024 banking and capital markets outlook*. <https://www.deloitte.com/us/en/insights/industry/financial-services/financial-services-industry-outlooks/banking-industry-outlook-2024.html>

- Edelman. (2024). *Edelman Trust Barometer 2024: Financial services report*. <https://www.edelmansmithfield.com/2024-edelman-trust-barometer/supplemental-report-insights-financial-service> Edelman Smithfield.
- Harvard Business Review. (2026, March 10). *The hidden causes of AI workslop—and how to fix them* [Audio podcast episode]. In *HBR IdeaCast*. <https://hbr.org/podcast/2026/03/the-hidden-causes-of-ai-workslop-and-how-to-fix-them>
- Lepczyk, A. (2024, June 24). *Share draft and auto loan penetration jump substantially*. *CreditUnions.com*. <https://creditunions.com/blogs/industry-insights/share-draft-and-auto-loan-penetration-jump-substantially>
- McKinsey & Company. (2023, June 14). *The economic potential of generative AI: The next productivity frontier*. [https://www.mckinsey.de/~media/mckinsey/locations/europe and middle east/deutschland/news/presse/2023/2023-06-14 mgi genai report 23/the-economic-potential-of-generative-ai-the-next-productivity-frontier-vf.pdf](https://www.mckinsey.de/~media/mckinsey/locations/europe%20and%20middle%20east/deutschland/news/presse/2023/2023-06-14%20mgi%20genai%20report%2023/the-economic-potential-of-generative-ai-the-next-productivity-frontier-vf.pdf)
- McKinsey & Company. (2025). *Superagency in the workplace: Empowering people to unlock AI's full potential*. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work>
- National Association of Realtors. (2024). *2024 profile of home buyers and sellers*. https://www.nar.realtor/sites/default/files/2024-11/2024-profile-of-home-buyers-and-sellers-highlights-11-04-2024_2.pdf
- National Credit Union Administration. (2023). *2022 annual report*. <https://ncua.gov/files/annual-reports/annual-report-2022.pdf>
- Oregon Community Credit Union. (2023). *2023 annual report*. <https://myoccu.org/about-occu/annual-report/2023-annual-report>
- Oregon Community Credit Union. (2024). *2024 annual report*. <https://myoccu.org/about-occu/annual-report/2024-annual-report>
- Oregon Community Credit Union. (n.d.). *About OCCU*. <https://myoccu.org/about-occu>
- Pew Research Center. (2023). *Financial attitudes and behaviors of younger adults*. <https://www.pewresearch.org>
- Pickens, J. (2026, February 3). *Interview conducted with Amy Durst and WCMS Project II team as part of WCMS Project II research* [Transcript of an unpublished interview]. Rogue Credit Union.
- Rogue Credit Union. (2025). *Strategic drivers and enablers* [Internal strategic planning document].
- Rogue Credit Union. (2026). *Leadership council strategic priorities* [Internal presentation].

- Shevlin, R. (2024). *What's going on in banking 2024: Finding the next wave to ride*. Cornerstone Advisors. [https://www.cornstone.com/hubfs/WGOIB 2024/2024-Whats-Going-On-In-Banking_Cornerstone-Advisors.pdf](https://www.cornstone.com/hubfs/WGOIB%2024/2024-Whats-Going-On-In-Banking_Cornerstone-Advisors.pdf)
- Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (2025, November 5). *The state of AI in 2025: Agents, innovation, and transformation*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Smith, G. (2024, June 27). *AI is still a delusion*. *Mind Matters*. <https://mindmatters.ai/2024/06/ai-is-still-a-delusion>
- Stephenson, M. (2026, February 5). *Interview conducted with Amy Durst and WCMS Project II team as part of WCMS Project II research* [Transcript of an unpublished interview]. Rogue Credit Union.
- U.S. Census Bureau. (2024). *Annual state resident population estimates: April 1, 2020 to July 1, 2023*. <https://www.census.gov/data/tables/time-series/demo/popest/2020s-state-total.html>

XII. APPENDIX

The economic potential of generative AI

The next productivity frontier

June 2023



Authors

Michael Chui
Eric Hazan
Roger Roberts
Alex Singla
Kate Smaje
Alex Sukharevsky
Lareina Yee
Rodney Zemmel

```

        modifier_ob.modifiers.new("mirror_ob")
        mirror_ob.mirror_object = mirror_ob

        operation == "MIRROR_X":
            mirror_mod.use_x = True
            mirror_mod.use_y = False
            mirror_mod.use_z = False
        operation == "MIRROR_Y":
            mirror_mod.use_x = False
            mirror_mod.use_y = True
            mirror_mod.use_z = False
        operation == "MIRROR_Z":
            mirror_mod.use_x = False
            mirror_mod.use_y = False
            mirror_mod.use_z = True

```

```

        selection at the end -add back the deselected objects
        mirror_ob.select= 1
        modifier_ob.select=1
        key.context.scene.objects.active = modifier_ob
        print("selected" + str(modifier_ob)) # modifier object selected
        mirror_ob.select = 0
        key.context.selected_objects[0]
        scene.objects[one.name].select = 1

```

```

print("please select exactly two objects,")

```

OPERATOR CLASSES

```

class Operator:
    def mirror_to_the_selected_object():
        mirror_ob = scene.objects[one.name]
        mirror_ob.mirror_mirror_x = True
        mirror_ob.mirror_mirror_y = True
        mirror_ob.mirror_mirror_z = True
        mirror_ob.mirror_x = True
        mirror_ob.mirror_y = True
        mirror_ob.mirror_z = True
        mirror_ob.mirror_object = mirror_ob
        mirror_ob.select = 1
        mirror_ob.name = "Mirror to the selected object"
        mirror_ob.active_object = None

```

Contents

Key insights

3

Chapter 1: Generative AI
as a technology catalyst

4

Glossary

6

Chapter 2: Generative AI use
cases across functions and
industries

8

Spotlight: Retail and
consumer packaged goods

27

Spotlight: Banking

28

Spotlight: Pharmaceuticals
and medical products

30

Chapter 3: The generative
AI future of work: Impacts
on work activities, economic
growth, and productivity


32

Chapter 4: Considerations
for businesses and society

48

Appendix

53



Key insights

1. *Generative AI's impact on productivity could add trillions of dollars in value to the global economy.* Our latest research estimates that generative AI could add the equivalent of \$2.6 trillion to \$4.4 trillion annually across the 63 use cases we analyzed—by comparison, the United Kingdom's entire GDP in 2021 was \$3.1 trillion. This would increase the impact of all artificial intelligence by 15 to 40 percent. This estimate would roughly double if we include the impact of embedding generative AI into software that is currently used for other tasks beyond those use cases.
2. *About 75 percent of the value that generative AI use cases could deliver falls across four areas: Customer operations, marketing and sales, software engineering, and R&D.* Across 16 business functions, we examined 63 use cases in which the technology can address specific business challenges in ways that produce one or more measurable outcomes. Examples include generative AI's ability to support interactions with customers, generate creative content for marketing and sales, and draft computer code based on natural-language prompts, among many other tasks.
3. *Generative AI will have a significant impact across all industry sectors.* Banking, high tech, and life sciences are among the industries that could see the biggest impact as a percentage of their revenues from generative AI. Across the banking industry, for example, the technology could deliver value equal to an additional \$200 billion to \$340 billion annually if the use cases were fully implemented. In retail and consumer packaged goods, the potential impact is also significant at \$400 billion to \$660 billion a year.
4. *Generative AI has the potential to change the anatomy of work, augmenting the capabilities of individual workers by automating some of their individual activities.* Current generative AI and other technologies have the potential to automate work activities that absorb 60 to 70 percent of employees' time today. In contrast, we previously estimated that technology has the potential to automate half of the time employees spend working.¹ The acceleration in the potential for technical automation is largely due to generative AI's increased ability to understand natural language, which is required for work activities that account for 25 percent of total work time. Thus, generative AI has more impact on knowledge work associated with occupations that have higher wages and educational requirements than on other types of work.
5. *The pace of workforce transformation is likely to accelerate, given increases in the potential for technical automation.* Our updated adoption scenarios, including technology development, economic feasibility, and diffusion timelines, lead to estimates that half of today's work activities could be automated between 2030 and 2060, with a midpoint in 2045, or roughly a decade earlier than in our previous estimates.
6. *Generative AI can substantially increase labor productivity across the economy, but that will require investments to support workers as they shift work activities or change jobs.* Generative AI could enable labor productivity growth of 0.1 to 0.6 percent annually through 2040, depending on the rate of technology adoption and redeployment of worker time into other activities. Combining generative AI with all other technologies, work automation could add 0.2 to 3.3 percentage points annually to productivity growth. However, workers will need support in learning new skills, and some will change occupations. If worker transitions and other risks can be managed, generative AI could contribute substantively to economic growth and support a more sustainable, inclusive world.
7. *The era of generative AI is just beginning.* Excitement over this technology is palpable, and early pilots are compelling. But a full realization of the technology's benefits will take time, and leaders in business and society still have considerable challenges to address. These include managing the risks inherent in generative AI, determining what new skills and capabilities the workforce will need, and rethinking core business processes such as retraining and developing new skills.



1

Generative AI as a technology catalyst

To grasp what lies ahead requires an understanding of the breakthroughs that have enabled the rise of generative AI, which were decades in the making. ChatGPT, GitHub Copilot, Stable Diffusion, and other generative AI tools that have captured current public attention are the result of significant levels of investment in recent years that have helped advance machine learning and deep learning. This investment undergirds the AI applications embedded in many of the products and services we use every day.

But because AI has permeated our lives incrementally—through everything from the tech powering our smartphones to autonomous-driving features on cars to the tools retailers use to surprise and delight consumers—its progress was almost imperceptible. Clear milestones, such as when AlphaGo, an AI-based program developed by DeepMind, defeated a world champion Go player in 2016, were celebrated but then quickly faded from the public's consciousness.

ChatGPT and its competitors have captured the imagination of people around the world in a way AlphaGo did not, thanks to their broad utility—almost anyone can use them to communicate and create—and preternatural ability to have a conversation with a user. The latest generative AI applications can perform a range of routine tasks, such as the reorganization and classification of data. But it is their ability to write text, compose music, and create digital art that has garnered headlines and persuaded consumers and households to experiment on their own. As a result, a broader set of stakeholders are grappling with generative AI's impact on business and society but without much context to help them make sense of it.

How did we get here? Gradually, then all of a sudden

For the purposes of this report, we define generative AI as applications typically built using foundation models. These models contain expansive artificial neural networks inspired by the billions of neurons connected in the human brain. Foundation models are part of what is called deep learning, a term that alludes to the many deep layers within neural networks. Deep learning has powered many of the recent advances in AI, but the foundation models powering generative AI applications are a step change evolution within deep learning. Unlike previous deep learning models, they can process extremely large and varied sets of unstructured data and perform more than one task.

Foundation models have enabled new capabilities and vastly improved existing ones across a broad range of modalities, including images, video, audio, and computer code. AI trained on these models can perform several functions; it can classify, edit, summarize, answer questions, and draft new content, among other tasks.

Continued innovation will also bring new challenges. For example, the computational power required to train generative AI with hundreds of billions of parameters threatens to become a bottleneck in development.² Further, there's a significant move—spearheaded by the open-source community and spreading to the leaders of generative AI companies themselves—to make AI more responsible, which could increase its costs.

Nonetheless, funding for generative AI, though still a fraction of total investments in artificial intelligence, is significant and growing rapidly—reaching a total of \$12 billion in the first five months of 2023 alone. Venture capital and other private external investments in generative AI increased by an average compound growth rate of 74 percent annually from 2017 to 2022. During the same period, investments in artificial intelligence overall rose annually by 29 percent, albeit from a higher base.

The rush to throw money at all things generative AI reflects how quickly its capabilities have developed. ChatGPT was released in November 2022. Four months later, OpenAI released a new large language model, or LLM, called GPT-4 with markedly improved capabilities.³ Similarly, by May 2023, Anthropic's generative AI, Claude, was able to process 100,000 tokens of text, equal to about 75,000 words in a minute—the length of the average novel—compared with roughly 9,000 tokens when it was introduced in March 2023.⁴ And in May 2023, Google announced several new features powered by generative AI, including Search Generative Experience and a new LLM called PaLM 2 that will power its Bard chatbot, among other Google products.⁵

From a geographic perspective, external private investment in generative AI, mostly from tech giants and venture capital firms, is largely concentrated in North America, reflecting the continent's current domination of the overall AI investment landscape. Generative AI-related companies based in the United States raised about \$8 billion from 2020 to 2022, accounting for 75 percent of total investments in such companies during that period.⁶

Generative AI has stunned and excited the world with its potential for reshaping how knowledge work gets done in industries and business functions across the entire economy. Across functions such as sales and marketing, customer operations, and software development, it is poised to transform roles and boost performance. In the process, it could unlock trillions of dollars in value across sectors from banking to life sciences. We have used two overlapping lenses in this report to understand the potential for generative AI to create value for companies and alter the workforce. The following sections share our initial findings.

Glossary

Application programming interface (API) is a way to programmatically access (usually external) models, data sets, or other pieces of software.

Artificial intelligence (AI) is the ability of software to perform tasks that traditionally require human intelligence.

Artificial neural networks (ANNs) are composed of interconnected layers of software-based calculators known as “neurons.” These networks can absorb vast amounts of input data and process that data through multiple layers that extract and learn the data’s features.

Deep learning is a subset of machine learning that uses deep neural networks, which are layers of connected “neurons” whose connections have parameters or weights that can be trained. It is especially effective at learning from unstructured data such as images, text, and audio.

Early and late scenarios are the extreme scenarios of our work-automation model. The “earliest” scenario flexes all parameters to the extremes of plausible assumptions, resulting in faster automation development and adoption, and the “latest” scenario flexes all parameters in the opposite direction. The reality is likely to fall somewhere between the two.

Fine-tuning is the process of adapting a pretrained foundation model to perform better in a specific task. This entails a relatively short period of training on a labeled data set, which is much smaller than the data set the model was initially trained on. This additional training allows the model to learn and adapt to the nuances, terminology, and specific patterns found in the smaller data set.

Foundation models (FM) are deep learning models trained on vast quantities of unstructured, unlabeled data that can be used for a wide range of tasks out of the box or adapted to specific tasks through fine-tuning. Examples of these models are GPT-4, PaLM, DALL-E 2, and Stable Diffusion.

Generative AI is AI that is typically built using foundation models and has capabilities that earlier AI did not have, such as the ability to generate content. Foundation models can also be used for nongenerative purposes (for example, classifying user sentiment as negative or positive based on call transcripts) while offering significant improvement over earlier models. For simplicity, when we refer to generative AI in this article, we include all foundation model use cases.

Graphics processing units (GPUs) are computer chips that were originally developed for producing computer graphics (such as for video games) and are also useful for deep learning applications. In contrast, traditional machine learning and other analyses usually run on central processing units (CPUs), normally referred to as a computer’s “processor.”

Large language models (LLMs) make up a class of foundation models that can process massive amounts of unstructured text and learn the relationships between words or portions of words, known as tokens. This enables LLMs to generate natural-language text, performing tasks such as summarization or knowledge extraction. GPT-4 (which underlies ChatGPT) and LaMDA (the model behind Bard) are examples of LLMs.

Machine learning (ML) is a subset of AI in which a model gains capabilities after it is trained on, or shown, many example data points. Machine learning algorithms detect patterns and learn how to make predictions and recommendations by processing data and experiences, rather than by receiving explicit programming instruction. The algorithms also adapt and can become more effective in response to new data and experiences.

Modality is a high-level data category such as numbers, text, images, video, and audio.

Productivity from labor is the ratio of GDP to total hours worked in the economy. Labor productivity growth comes from increases in the amount of capital available to each worker, the education and experience of the workforce, and improvements in technology.

Prompt engineering refers to the process of designing, refining, and optimizing input prompts to guide a generative AI model toward producing desired (that is, accurate) outputs.

Self-attention, sometimes called intra-attention, is a mechanism that aims to mimic cognitive attention, relating different positions of a single sequence to compute a representation of the sequence.


Structured data are tabular data (for example, organized in tables, databases, or spreadsheets) that can be used to train some machine learning models effectively.

Transformers are a relatively new neural network architecture that relies on self-attention mechanisms to transform a sequence of inputs into a sequence of outputs while focusing its attention on important parts of the context around the inputs. Transformers do not rely on convolutions or recurrent neural networks.

Technical automation potential refers to the share of the worktime that could be automated. We assessed the technical potential for automation across the global economy through an analysis of the component activities of each occupation. We used databases published by institutions including the World Bank and the US Bureau of Labor Statistics to break down about 850 occupations into approximately 2,100 activities, and we determined the performance capabilities needed for each activity based on how humans currently perform them.

Use cases are targeted applications to a specific business challenge that produces one or more measurable outcomes. For example, in marketing, generative AI could be used to generate creative content such as personalized emails.

Unstructured data lack a consistent format or structure (for example, text, images, and audio files) and typically require more advanced techniques to extract insights.




Generative AI use cases across functions and industries

Generative AI is a step change in the evolution of artificial intelligence. As companies rush to adapt and implement it, understanding the technology's potential to deliver value to the economy and society at large will help shape critical decisions. We have used two complementary lenses to determine where generative AI with its current capabilities could deliver the biggest value and how big that value could be (Exhibit 1).

Exhibit 1

The potential impact of generative AI can be evaluated through two lenses.



¹For quantitative analysis, revenue impacts were recast as productivity increases on the corresponding spend in order to maintain comparability with cost impacts and not to assume additional growth in any particular market.

McKinsey & Company

The first lens scans use cases for generative AI that organizations could adopt. We define a “use case” as a targeted application of generative AI to a specific business challenge, resulting in one or more measurable outcomes. For example, a use case in marketing is the application of generative AI to generate creative content such as personalized emails, the measurable outcomes of which potentially include reductions in the cost of generating such content and increases in revenue from the enhanced effectiveness of higher-quality content at scale. We identified 63 generative AI use cases spanning 16 business functions that could deliver total value in the range of \$2.6 trillion to \$4.4 trillion in economic benefits annually when applied across industries.

That would add 15 to 40 percent to the \$11.0 trillion to \$17.7 trillion of economic value that we now estimate nongenerative artificial intelligence and analytics could unlock. (Our previous estimate from 2017 was that AI could deliver \$9.5 trillion to \$15.4 trillion in economic value.)


Our second lens complements the first by analyzing generative AI’s potential impact on the work activities required in some 850 occupations. We modeled scenarios to estimate when generative AI could perform each of more than 2,100 “detailed work activities”—such as “communicating with others about operational plans or activities”—that make up those occupations across the world economy. This enables us to estimate how the current capabilities of generative AI could affect labor productivity across all work currently done by the global workforce.

Some of this impact will overlap with cost reductions in the use case analysis described above, which we assume are the result of improved labor productivity. Netting out this overlap, the total economic benefits of generative AI—including the major use cases we explored and the myriad increases in productivity that are likely to materialize when the technology is applied across knowledge workers' activities—amounts to \$6.1 trillion to \$7.9 trillion annually (Exhibit 2).

Exhibit 2

Generative AI could create additional value potential above what could be unlocked by other AI and analytics.

AI's potential impact on the global economy, \$ trillion



¹Updated use case estimates from "Notes from the AI frontier: Applications and value of deep learning," McKinsey Global Institute, April 17, 2018.

McKinsey & Company

While generative AI is an exciting and rapidly advancing technology, the other applications of AI discussed in our previous report continue to account for the majority of the overall potential value of AI. Traditional advanced-analytics and machine learning algorithms are highly effective at performing numerical and optimization tasks such as predictive modeling, and they continue to find new applications in a wide range of industries. However, as generative AI continues to develop and mature, it has the potential to open wholly new frontiers in creativity and innovation. It has already expanded the possibilities of what AI overall can achieve (please see Box 1, “How we estimated the value potential of generative AI use cases”).

Box 1

How we estimated the value potential of generative AI use cases

To assess the potential value of generative AI, we updated a proprietary McKinsey database of potential AI use cases and drew on the experience of more than 100 experts in industries and their business functions.¹ Our updates examined use cases of generative AI—specifically, how generative AI techniques (primarily transformer-based neural networks) can be used to solve problems not well addressed by previous technologies.

We analyzed only use cases for which generative AI could deliver a significant improvement in the outputs that drive key value. In particular, our estimates of the primary value the technology could unlock do not include use cases for which the sole benefit would be its ability to use natural language. For example, natural-language capabilities would be the key driver of value in

a customer service use case but not in a use case optimizing a logistics network, where value primarily arises from quantitative analysis.

We then estimated the potential annual value of these generative AI use cases if they were adopted across the entire economy. For use cases aimed at increasing revenue, such as some of those in sales and marketing, we estimated the economy-wide value generative AI could deliver by increasing the productivity of sales and marketing expenditures.

Our estimates are based on the structure of the global economy in 2022 and do not consider the value generative AI could create if it produced entirely new product or service categories.

¹ “Notes from the AI frontier: Applications and value of deep learning,” McKinsey Global Institute, April 17, 2018.


In this chapter, we highlight the value potential of generative AI across two dimensions: business function and modality.

Value potential by function

While generative AI could have an impact on most business functions, a few stand out when measured by the technology's impact as a share of functional cost (Exhibit 3). Our analysis of 16 business functions identified just four—customer operations, marketing and sales, software engineering, and research and development—that could account for approximately 75 percent of the total annual value from generative AI use cases.

Exhibit 3

Using generative AI in just a few functions could drive most of the technology's impact across potential corporate use cases.



Note: Impact is averaged.

¹Excluding software engineering.

Source: Comparative Industry Service (CIS), IHS Markit; Oxford Economics; McKinsey Corporate and Business Functions database; McKinsey Manufacturing and Supply Chain 360; McKinsey Sales Navigator; Ignite, a McKinsey database; McKinsey analysis

McKinsey & Company

Notably, the potential value of using generative AI for several functions that were prominent in our previous sizing of AI use cases, including manufacturing and supply chain functions, is now much lower.⁷ This is largely explained by the nature of generative AI use cases, which exclude most of the numerical and optimization applications that were the main value drivers for previous applications of AI.

Generative AI as a virtual expert

In addition to the potential value generative AI can deliver in function-specific use cases, the technology could drive value across an entire organization by revolutionizing internal knowledge management systems. Generative AI's impressive command of natural-language processing can help employees retrieve stored internal knowledge by formulating queries in the same way they might ask a human a question and engage in continuing dialogue. This could empower teams to quickly access relevant information, enabling them to rapidly make better-informed decisions and develop effective strategies.

In 2012, the McKinsey Global Institute (MGI) estimated that knowledge workers spent about a fifth of their time, or one day each work week, searching for and gathering information. If generative AI could take on such tasks, increasing the efficiency and effectiveness of the workers doing them, the benefits would be huge. Such virtual expertise could rapidly “read” vast libraries of corporate information stored in natural language and quickly scan source material in dialogue with a human who helps fine-tune and tailor its research, a more scalable solution than hiring a team of human experts for the task.

Following are examples of how generative AI could produce operational benefits as a virtual expert in a handful of use cases.

In addition to the potential value generative AI can deliver in specific use cases, the technology could drive value across an entire organization by revolutionizing internal knowledge management systems.

How customer operations could be transformed




Customer self-service interactions

Customer interacts with a humanlike chatbot that delivers immediate, personalized responses to complex inquiries, ensuring a consistent brand voice regardless of customer language or location.

Customer-agent interactions

Human agent uses AI-developed call scripts and receives real-time assistance and suggestions for responses during phone conversations, instantly accessing relevant customer data for tailored and real-time information delivery.



Agent self-improvement

Agent receives a summarization of the conversation in a few succinct points to create a record of customer complaints and actions taken.

Agent uses automated, personalized insights generated by AI, including tailored follow-up messages or personalized coaching suggestions.

Customer operations

Generative AI has the potential to revolutionize the entire customer operations function, improving the customer experience and agent productivity through digital self-service and enhancing and augmenting agent skills. The technology has already gained traction in customer service because of its ability to automate interactions with customers using natural language. Research found that at one company with 5,000 customer service agents, the application of generative AI increased issue resolution by 14 percent an hour and reduced the time spent handling an issue by 9 percent.⁸ It also reduced agent attrition and requests to speak to a manager by 25 percent. Crucially, productivity and quality of service improved most among less-experienced agents, while the AI assistant did not increase—and sometimes decreased—the productivity and quality metrics of more highly skilled agents. This is because AI assistance helped less-experienced agents communicate using techniques similar to those of their higher-skilled counterparts.

The following are examples of the operational improvements generative AI can have for specific use cases:

- **Customer self-service.** Generative AI–fueled chatbots can give immediate and personalized responses to complex customer inquiries regardless of the language or location of the customer. By improving the quality and effectiveness of interactions via automated channels, generative AI could automate responses to a higher percentage of customer inquiries, enabling customer care teams to take on inquiries that can only be resolved by a human agent. Our research found that roughly half of customer contacts made by banking, telecommunications, and utilities companies in North America are already handled by machines, including but not exclusively AI. We estimate that generative AI could further reduce the volume of human-serviced contacts by up to 50 percent, depending on a company's existing level of automation.
- **Resolution during initial contact.** Generative AI can instantly retrieve data a company has on a specific customer, which can help a human customer service representative more successfully answer questions and resolve issues during an initial interaction.
- **Reduced response time.** Generative AI can cut the time a human sales representative spends responding to a customer by providing assistance in real time and recommending next steps.
- **Increased sales.** Because of its ability to rapidly process data on customers and their browsing histories, the technology can identify product suggestions and deals tailored to customer preferences. Additionally, generative AI can enhance quality assurance and coaching by gathering insights from customer conversations, determining what could be done better, and coaching agents.

We estimate that applying generative AI to customer care functions could increase productivity at a value ranging from 30 to 45 percent of current function costs.

Our analysis captures only the direct impact generative AI might have on the productivity of customer operations. It does not account for potential knock-on effects the technology may have on customer satisfaction and retention arising from an improved experience, including better understanding of the customer's context that can assist human agents in providing more personalized help and recommendations.

How marketing and sales could be transformed



Strategization

Sales and marketing professionals efficiently gather market trends and customer information from unstructured data sources (for example, social media, news, research, product information, and customer feedback) and draft effective marketing and sales communications.


Awareness

Customers see campaigns tailored to their segment, language, and demographic.



Consideration

Customers can access comprehensive information, comparisons, and dynamic recommendations, such as personal “try ons.”



Retention

Customers are more likely to be retained with customized messages and rewards, and they can interact with AI-powered customer-support chatbots that manage the relationship proactively, with fewer escalations to human agents.

Conversion

Virtual sales representatives enabled by generative AI emulate humanlike qualities—such as empathy, personalized communication, and natural language processing—to build trust and rapport with customers.



Marketing and sales

Generative AI has taken hold rapidly in marketing and sales functions, in which text-based communications and personalization at scale are driving forces. The technology can create personalized messages tailored to individual customer interests, preferences, and behaviors, as well as do tasks such as producing first drafts of brand advertising, headlines, slogans, social media posts, and product descriptions.

However, introducing generative AI to marketing functions requires careful consideration. For one thing, mathematical models trained on publicly available data without sufficient safeguards against plagiarism, copyright violations, and branding recognition risks infringing on intellectual property rights. A virtual try-on application may produce biased representations of certain demographics because of limited or biased training data. Thus, significant human oversight is required for conceptual and strategic thinking specific to each company's needs.

Potential operational benefits from using generative AI for marketing include the following:

- **Efficient and effective content creation.** Generative AI could significantly reduce the time required for ideation and content drafting, saving valuable time and effort. It can also facilitate consistency across different pieces of content, ensuring a uniform brand voice, writing style, and format. Team members can collaborate via generative AI, which can integrate their ideas into a single cohesive piece. This would allow teams to significantly enhance personalization of marketing messages aimed at different customer segments, geographies, and demographics. Mass email campaigns can be instantly translated into as many languages as needed, with different imagery and messaging depending on the audience. Generative AI's ability to produce content with varying specifications could increase customer value, attraction, conversion, and retention over a lifetime and at a scale beyond what is currently possible through traditional techniques.
- **Enhanced use of data.** Generative AI could help marketing functions overcome the challenges of unstructured, inconsistent, and disconnected data—for example, from different databases—by interpreting abstract data sources such as text, image, and varying structures. It can help marketers better use data such as territory performance, synthesized customer feedback, and customer behavior to generate data-informed marketing strategies such as targeted customer profiles and channel recommendations. Such tools could identify and synthesize trends, key drivers, and market and product opportunities from unstructured data such as social media, news, academic research, and customer feedback.
- **SEO optimization.** Generative AI can help marketers achieve higher conversion and lower cost through search engine optimization (SEO) for marketing and sales technical components such as page titles, image tags, and URLs. It can synthesize key SEO tokens, support specialists in SEO digital content creation, and distribute targeted content to customers.
- **Product discovery and search personalization.** With generative AI, product discovery and search can be personalized with multimodal inputs from text, images and speech, and deep understanding of customer profiles. For example, technology can leverage individual user preferences, behavior, and purchase history to help customers discover the most relevant products and generate personalized product descriptions. This would allow CPG, travel, and retail companies to improve their ecommerce sales by achieving higher website conversion rates.

We estimate that generative AI could increase the productivity of the marketing function with a value between 5 and 15 percent of total marketing spending.

Our analysis of the potential use of generative AI in marketing doesn't account for knock-on effects beyond the direct impacts on productivity. Generative AI-enabled synthesis could provide higher-quality data insights, leading to new ideas for marketing campaigns and better-targeted customer segments. Marketing functions could shift resources to producing higher-quality content for owned channels, potentially reducing spending on external channels and agencies.

Generative AI could also change the way both B2B and B2C companies approach sales. The following are two use cases for sales:

- **Increase probability of sale.** Generative AI could identify and prioritize sales leads by creating comprehensive consumer profiles from structured and unstructured data and suggesting actions to staff to improve client engagement at every point of contact. For example, generative AI could provide better information about client preferences, potentially improving close rates.
- **Improve lead development.** Generative AI could help sales representatives nurture leads by synthesizing relevant product sales information and customer profiles and creating discussion scripts to facilitate customer conversation, including up- and cross-selling talking points. It could also automate sales follow-ups and passively nurture leads until clients are ready for direct interaction with a human sales agent.

Our analysis suggests that implementing generative AI could increase sales productivity by approximately 3 to 5 percent of current global sales expenditures.

This analysis may not fully account for additional revenue that generative AI could bring to sales functions. For instance, generative AI's ability to identify leads and follow-up capabilities could uncover new leads and facilitate more effective outreach that would bring in additional revenue. Also, the time saved by sales representatives due to generative AI's capabilities could be invested in higher-quality customer interactions, resulting in increased sales success.

Generative AI as a virtual collaborator

In other cases, generative AI can drive value by working in partnership with workers, augmenting their work in ways that accelerate their productivity. Its ability to rapidly digest mountains of data and draw conclusions from it enables the technology to offer insights and options that can dramatically enhance knowledge work. This can significantly speed up the process of developing a product and allow employees to devote more time to higher-impact tasks.

Generative AI could increase sales productivity by 3 to 5 percent of current global sales expenditures.

How software engineering could be transformed



Inception and planning

Software engineers and product managers use generative AI to assist in analyzing, cleaning, and labeling large volumes of data, such as user feedback, market trends, and existing system logs.

System design

Engineers use generative AI to create multiple IT architecture designs and iterate on the potential configurations, accelerating system design, and allowing faster time to market.



Coding

Engineers are assisted by AI tools that can code, reducing development time by assisting with drafts, rapidly finding prompts, and serving as an easily navigable knowledge base.

Testing

Engineers employ algorithms that can enhance functional and performance testing to ensure quality and can generate test cases and test data automatically.





Maintenance

Engineers use AI insights on system logs, user feedback, and performance data to help diagnose issues, suggest fixes, and predict other high-priority areas of improvement.

Software engineering

Treating computer languages as just another language opens new possibilities for software engineering. Software engineers can use generative AI in pair programming and to do augmented coding and train LLMs to develop applications that generate code when given a natural-language prompt describing what that code should do.

Software engineering is a significant function in most companies, and it continues to grow as all large companies, not just tech titans, embed software in a wide array of products and services. For example, much of the value of new vehicles comes from digital features such as adaptive cruise control, parking assistance, and IoT connectivity.

According to our analysis, the direct impact of AI on the productivity of software engineering could range from 20 to 45 percent of current annual spending on the function. This value would arise primarily from reducing time spent on certain activities, such as generating initial code drafts, code correction and refactoring, root-cause analysis, and generating new system designs. By accelerating the coding process, generative AI could push the skill sets and capabilities needed in software engineering toward code and architecture design. One study found that software developers using Microsoft's GitHub Copilot completed tasks 56 percent faster than those not using the tool.⁹ An internal McKinsey empirical study of software engineering teams found those who were trained to use generative AI tools rapidly reduced the time needed to generate and refactor code—and engineers also reported a better work experience, citing improvements in happiness, flow, and fulfillment.

Our analysis did not account for the increase in application quality and the resulting boost in productivity that generative AI could bring by improving code or enhancing IT architecture—which can improve productivity across the IT value chain. However, the quality of IT architecture still largely depends on software architects, rather than on initial drafts that generative AI's current capabilities allow it to produce.

Large technology companies are already selling generative AI for software engineering, including GitHub Copilot, which is now integrated with OpenAI's GPT-4, and Replit, used by more than 20 million coders.¹⁰

How product R&D could be transformed



Early research analysis

Researchers use generative AI to enhance market reporting, ideation, and product or solution drafting.



Virtual design

Researchers use generative AI to generate prompt-based drafts and designs, allowing them to iterate quickly with more design options.



Virtual simulations

Researchers accelerate and optimize the virtual simulation phase if combined with new deep learning generative design techniques.



Physical test planning

Researchers optimize test cases for more efficient testing, reducing the time required for physical build and testing.

Product R&D

Generative AI's potential in R&D is perhaps less well recognized than its potential in other business functions. Still, our research indicates the technology could deliver productivity with a value ranging from 10 to 15 percent of overall R&D costs.

For example, the life sciences and chemical industries have begun using generative AI foundation models in their R&D for what is known as generative design. Foundation models can generate candidate molecules, accelerating the process of developing new drugs and materials. Entos, a biotech pharmaceutical company, has paired generative AI with automated synthetic development tools to design small-molecule therapeutics. But the same principles can be applied to the design of many other products, including larger-scale physical products and electrical circuits, among others.

While other generative design techniques have already unlocked some of the potential to apply AI in R&D, their cost and data requirements, such as the use of “traditional” machine learning, can limit their application. Pretrained foundation models that underpin generative AI, or models that have been enhanced with fine-tuning, have much broader areas of application than models optimized for a single task. They can therefore accelerate time to market and broaden the types of products to which generative design can be applied. For now, however, foundation models lack the capabilities to help design products across all industries.

In addition to the productivity gains that result from being able to quickly produce candidate designs, generative design can also enable improvements in the designs themselves, as in the following examples of the operational improvements generative AI could bring:

- **Enhanced design.** Generative AI can help product designers reduce costs by selecting and using materials more efficiently. It can also optimize designs for manufacturing, which can lead to cost reductions in logistics and production.
- **Improved product testing and quality.** Using generative AI in generative design can produce a higher-quality product, resulting in increased attractiveness and market appeal. Generative AI can help to reduce testing time of complex systems and accelerate trial phases involving customer testing through its ability to draft scenarios and profile testing candidates.

We also identified a new R&D use case for nongenerative AI: deep learning surrogates, the use of which has grown since our earlier research, can be paired with generative AI to produce even greater benefits (see Box 2, “Deep learning surrogates”). To be sure, integration will require the development of specific solutions, but the value could be significant because deep learning surrogates have the potential to accelerate the testing of designs proposed by generative AI.

While we have estimated the potential direct impacts of generative AI on the R&D function, we did not attempt to estimate the technology’s potential to create entirely novel product categories. These are the types of innovations that can produce step changes not only in the performance of individual companies but in economic growth overall.

Box 2

Deep learning surrogates

Product design in industries producing physical products often involves physics-based virtual simulations such as computational fluid dynamics (CFD) and finite element analysis (FEA). Although they are faster than actual physical testing, these techniques can be time- and resource-intensive, especially for designing complex parts—running CFD simulations on graphics processing units

can take hours. And these techniques are even more complex and compute-intensive when they involve simulations coupled across multiple disciplines (for example, physical stress and temperature distribution), which is sometimes called multiphysics.

Deep learning applications are now revolutionizing the virtual testing phase of

the R&D process by using deep learning models to emulate (multi)physics-based simulations at higher speeds and lower costs. Instead of taking hours to run physics-based models, these deep learning surrogates can produce the results of simulations in just a few seconds, allowing researchers to test many more designs and enabling faster decision making on products and designs.

Value potential by modality

Technology has revolutionized the way we conduct business, and text-based AI is on the frontier of this change. Indeed, text-based data is plentiful, accessible, and easily processed and analyzed at large scale by LLMs, which has prompted a strong emphasis on them in the initial stages of generative AI development. The current investment landscape in generative AI is also heavily focused on text-based applications such as chatbots, virtual assistants, and language translation. However, we estimate that almost one-fifth of the value that generative AI can unlock across our use cases would take advantage of multimodal capabilities beyond text to text.

While most of generative AI's initial traction has been in text-based use cases, recent advances in generative AI have also led to breakthroughs in image generation, as OpenAI's DALL-E and Stable Diffusion have so amply illustrated, and much progress is being made in audio, including voice and music, and video. These capabilities have obvious applications in marketing for generating advertising materials and other marketing content, and these technologies are already being applied in media industries, including game design. Indeed, some of these examples challenge existing business models around talent, monetization, and intellectual property.¹¹

The multimodal capabilities of generative AI could also be used effectively in R&D. Generative AI systems could create first drafts of circuit designs, architectural drawings, structural engineering designs, and thermal designs based on prompts that describe requirements for a product. Achieving this will require training foundation models in these domains (think of LLMs trained on "design languages"). Once trained, such foundation models could increase productivity on a similar magnitude to software development.

Value potential by industry

Across the 63 use cases we analyzed, generative AI has the potential to generate \$2.6 trillion to \$4.4 trillion in value across industries. Its precise impact will depend on a variety of factors, such as the mix and importance of different functions, as well as the scale of an industry's revenue (Exhibit 4).

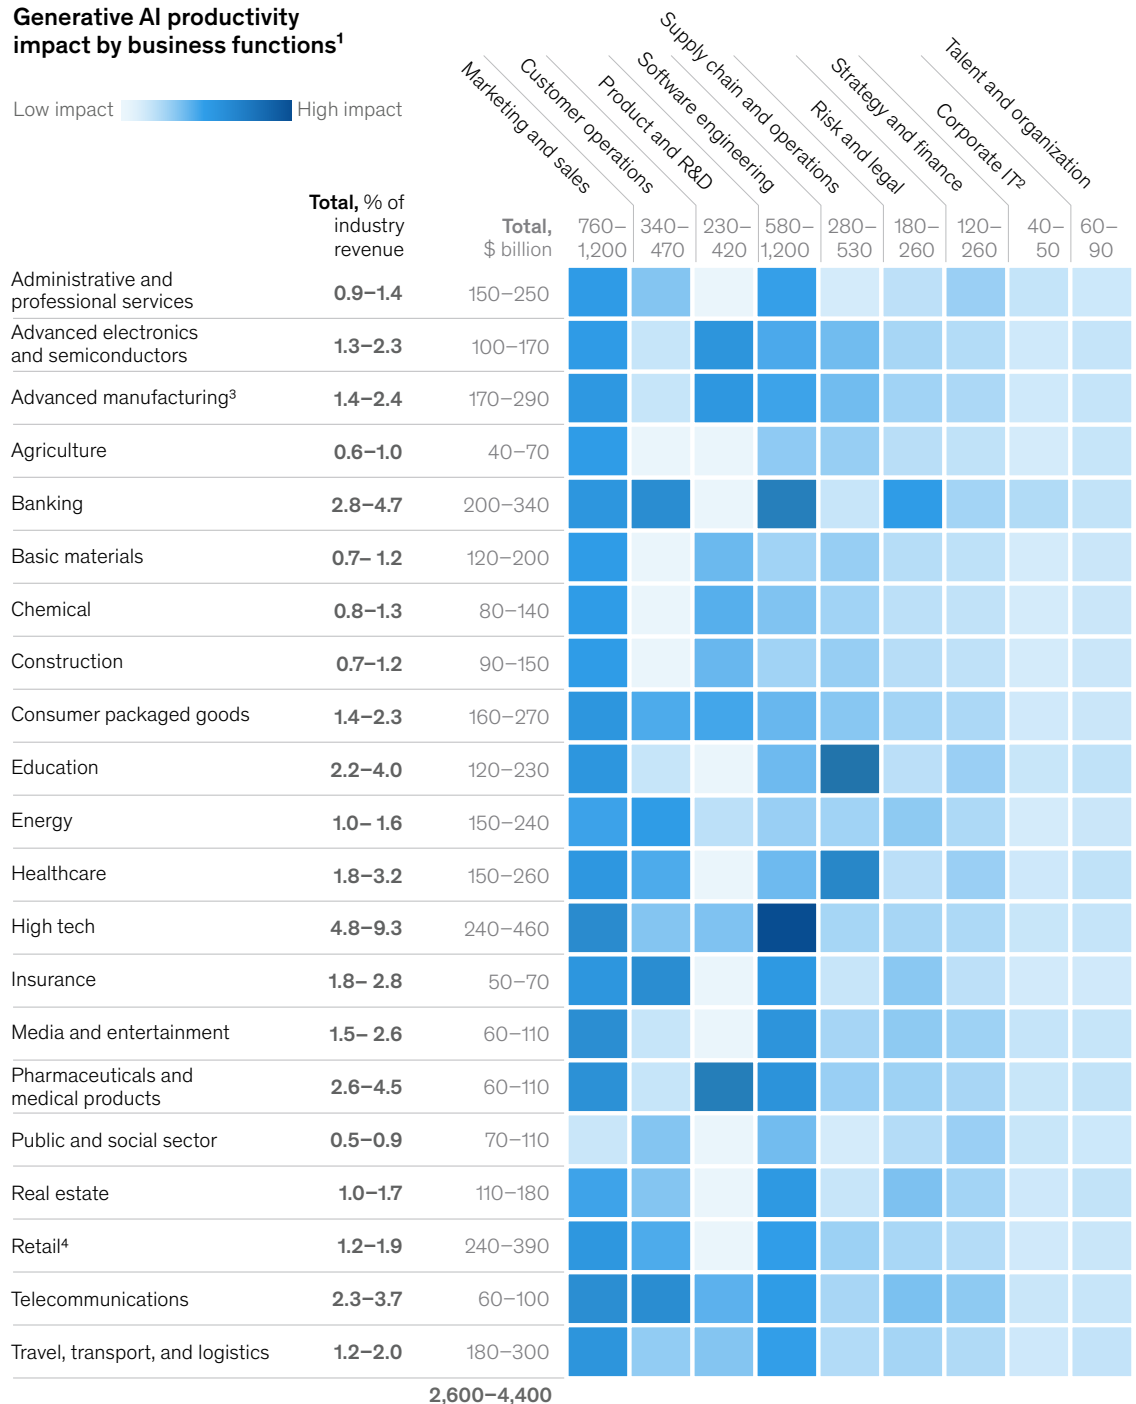
**Across 63 use cases,
generative AI has the
potential to generate
\$2.6 trillion to \$4.4 trillion
in value across industries.**

Exhibit 4

Generative AI use cases will have different impacts on business functions across industries.

Generative AI productivity impact by business functions¹

Low impact  High impact



Note: Figures may not sum to 100%, because of rounding.

¹Excludes implementation costs (eg, training, licenses).

²Excluding software engineering.

³Includes aerospace, defense, and auto manufacturing.

⁴Including auto retail.

Source: Comparative Industry Service (CIS), IHS Markit; Oxford Economics; McKinsey Corporate and Business Functions database; McKinsey Manufacturing and Supply Chain 360; McKinsey Sales Navigator; Ignite, a McKinsey database; McKinsey analysis

McKinsey & Company

For example, our analysis estimates generative AI could contribute roughly \$310 billion in additional value for the retail industry (including auto dealerships) by boosting performance in functions such as marketing and customer interactions. By comparison, the bulk of potential value in high tech comes from generative AI's ability to increase the speed and efficiency of software development (Exhibit 5).

Exhibit 5

Generative AI could deliver significant value when deployed in some use cases across a selection of top industries.

Selected examples of key use cases for main functional value drivers (nonexhaustive)

	Total value potential per industry, \$ billion (% of industry revenue)	Value potential, as % of operating profits ¹	Product R&D, software engineering	Customer operations	Marketing and sales	Other functions
Banking	200–340 (3–5%)	9–15	■ Legacy code conversion Optimize migration of legacy frameworks with natural-language translation capabilities	■ Customer emergency interactive voice response (IVR) Partially automate, accelerate, and enhance resolution rate of customer emergencies through generative AI-enhanced IVR interactions (eg, for credit card losses)	■ Custom retail banking offers Push personalized marketing and sales content tailored for each client of the bank based on profile and history (eg, personalized nudges), and generate alternatives for A/B testing	■ Risk model documentation Create model documentation, and scan for missing documentation and relevant regulatory updates
Retail and consumer packaged goods²	400–660 (1–2%)	27–44	■ Consumer research Accelerate consumer research by testing scenarios, and enhance customer targeting by creating “synthetic customers” to practice with	■ Augmented reality-assisted customer support Rapidly inform the workforce in real time about the status of products and consumer preferences	■ Assist copy writing for marketing content creation Accelerate writing of copy for marketing content and advertising scripts	■ Procurement suppliers process enhancement Draft playbooks for negotiating with suppliers
Pharma and medical products	60–110 (3–5%)	15–25	■ Research and drug discovery Accelerate the selection of proteins and molecules best suited as candidates for new drug formulation	■ Customer documentation generation Draft medication instructions and risk notices for drug resale	■ Generate content for commercial representatives Prepare scripts for interactions with physicians	■ Contract generation Draft legal documents incorporating specific regulatory requirements

¹Operating profit based on average profitability of selected industries in the 2020–22 period.

²Includes auto retail.

McKinsey & Company

In the banking industry, generative AI has the potential to improve on efficiencies already delivered by artificial intelligence by taking on lower-value tasks in risk management, such as required reporting, monitoring regulatory developments, and collecting data. In the life sciences industry, generative AI is poised to make significant contributions to drug discovery and development.

We share our detailed analysis of these industries in the following industry spotlights.

Generative AI could change the game for retail and consumer packaged goods companies

The technology could generate value for the retail and consumer packaged goods (CPG) industry by increasing productivity by 1.2 to 2.0 percent of annual revenues, or an additional \$400 billion to \$660 billion.¹ To streamline processes, generative AI could automate key functions such as customer service, marketing and sales, and inventory and supply chain management.

Technology has played an essential role in the retail and CPG industries for decades. Traditional AI and advanced-analytics solutions have helped companies manage vast pools of data across large numbers of SKUs, expansive supply chain and warehousing networks, and complex product categories such as consumables.

In addition, the industries are heavily customer facing, which offers opportunities for generative AI to complement previously existing artificial intelligence. For example, generative AI's ability to personalize offerings could optimize marketing and sales activities already handled by existing AI solutions. Similarly, generative AI tools excel at data management and could support existing AI-driven pricing tools. Applying generative AI to such activities could be a step toward integrating applications across a full enterprise.

Generative AI is already at work in some retail and CPG companies:

Reinvention of the customer interaction pattern

Consumers increasingly seek customization in everything from clothing and cosmetics to curated shopping experiences, personalized outreach, and food—and generative AI can improve that experience. Generative AI can aggregate market data to test concepts, ideas, and models. Stitch Fix, which uses algorithms to suggest style choices to its customers, has experimented with DALL·E to visualize products based on customer preferences regarding color, fabric, and

style. Using text-to-image generation, the company's stylists can visualize an article of clothing based on a consumer's preferences and then identify a similar article among Stitch Fix's inventory.

Retailers can create applications that give shoppers a next-generation experience, creating a significant competitive advantage in an era when customers expect to have a single natural-language interface help them select products. For example, generative AI can improve the process of choosing and ordering ingredients for a meal or preparing food—imagine a chatbot that could pull up the most popular tips from the comments attached to a recipe. There is also a big opportunity to enhance customer value management by delivering personalized marketing campaigns through a chatbot. Such applications can have human-like conversations about products in ways that can increase customer satisfaction, traffic, and brand loyalty. Generative AI offers retailers and CPG companies many opportunities to cross-sell and upsell, collect insights to improve product offerings, and increase their customer base, revenue opportunities, and overall marketing ROI.

Accelerating the creation of value in key areas

Generative AI tools can facilitate copy writing for marketing and sales, help brainstorm creative marketing ideas, expedite consumer research, and accelerate content analysis and creation. The potential improvement in writing and visuals can increase awareness and improve sales conversion rates.

Rapid resolution and enhanced insights in customer care

The growth of e-commerce also elevates the importance of effective consumer interactions. Retailers can combine existing AI tools with generative AI to enhance the capabilities of chatbots, enabling them to better mimic the interaction style of human agents—for

example, by responding directly to a customer's query, tracking or canceling an order, offering discounts, and upselling. Automating repetitive tasks allows human agents to devote more time to handling complicated customer problems and obtaining contextual information.

Disruptive and creative innovation

Generative AI tools can enhance the process of developing new versions of products by digitally creating new designs rapidly. A designer can generate packaging designs from scratch or generate variations on an existing design. This technology is developing rapidly and has the potential to add text-to-video generation.

Additional factors to consider

As retail and CPG executives explore how to integrate generative AI in their operations, they should keep in mind several factors that could affect their ability to capture value from the technology.

External inference. Generative AI has increased the need to understand whether generated content is based on fact or inference, requiring a new level of quality control.

Adversarial attacks. Foundation models are a prime target for attack by hackers and other bad actors, increasing the variety of potential security vulnerabilities and privacy risks.

To address these concerns, retail and CPG companies will need to strategically keep humans in the loop and ensure security and privacy are top considerations for any implementation. Companies will need to institute new quality checks for processes previously handled by humans, such as emails written by customer reps, and perform more-detailed quality checks on AI-assisted processes such as product design.

¹ Vehicular retail is included as part of our overall retail analysis.

Banks could realize substantial value from generative AI

Generative AI could have a significant impact on the banking industry, generating value from increased productivity of 2.8 to 4.7 percent of the industry's annual revenues, or an additional \$200 billion to \$340 billion. On top of that impact, the use of generative AI tools could also enhance customer satisfaction, improve decision making and employee experience, and decrease risks through better monitoring of fraud and risk.

Banking, a knowledge and technology-enabled industry, has already benefited significantly from previously existing applications of artificial intelligence in areas such as marketing and customer operations.¹ Generative AI applications could deliver additional benefits, especially because text modalities are prevalent in areas such as regulations and programming language, and the industry is customer facing, with many B2C and small-business customers.²

Several characteristics position the industry for the integration of generative AI applications:

- *Sustained digitization efforts along with legacy IT systems.* Banks have been investing in technology for decades, accumulating a significant amount of technical debt along with a siloed and complex IT architecture.³
- *Large customer-facing workforces.* Banking relies on a large number of service representatives such as call-center agents and wealth management financial advisers.
- *A stringent regulatory environment.* As a heavily regulated industry,

banking has a substantial number of risk, compliance, and legal needs.

- *White-collar industry.* Generative AI's impact could span the organization, assisting all employees in writing emails, creating business presentations, and other tasks.

On the move

Banks have started to grasp the potential of generative AI in their front lines and in their software activities. Early adopters are harnessing solutions such as ChatGPT as well as industry-specific solutions, primarily for software and knowledge applications. Three uses demonstrate its value potential to the industry:

A virtual expert to augment employee performance

A generative AI bot trained on proprietary knowledge such as policies, research, and customer interaction could provide always-on, deep technical support. Today, frontline spending is dedicated mostly to validating offers and interacting with clients, but giving frontline workers access to data as well could improve the customer experience. The technology could also monitor industries and clients and send alerts on semantic queries from public sources. For example, Morgan Stanley is building an AI assistant using GPT-4, with the aim of helping tens of thousands of wealth managers quickly find and synthesize answers from a massive internal knowledge base.⁴ The model combines search and content creation so wealth managers can find and tailor information for any client at any moment.

One European bank has leveraged generative AI to develop an environmental, social, and governance (ESG) virtual expert by synthesizing and extracting from long documents with unstructured information. The model answers complex questions based on a prompt, identifying the source of each answer and extracting information from pictures and tables.

Generative AI could reduce the significant costs associated with back-office operations. Such customer-facing chatbots could assess user requests and select the best service expert to address them based on characteristics such as topic, level of difficulty, and type of customer. Through generative AI assistants, service professionals could rapidly access all relevant information such as product guides and policies to instantaneously address customer requests.

Code acceleration to reduce tech debt and deliver software faster

Generative AI tools are useful for software development in four broad categories. First, they can draft code based on context via input code or natural language, helping developers code more quickly and with reduced friction while enabling automatic translations and no- and low-code tools. Second, such tools can automatically generate, prioritize, run, and review different code tests, accelerating testing and increasing coverage and effectiveness. Third, generative AI's natural-language translation capabilities can optimize the integration and migration of legacy frameworks. Last, the tools can review code to identify defects and inefficiencies in computing. The result is more robust, effective code.

¹ "Building the AI bank of the future," McKinsey, May 2021.

² McKinsey's Global Banking Annual Review, December 1, 2022.

³ Akhil Babbar, Raghavan Janardhanan, Remy Paternoster, and Henning Soller, "Why most digital banking transformations fail—and how to flip the odds," McKinsey, April 11, 2023.

⁴ Hugh Son, "Morgan Stanley is testing an OpenAI-powered chatbot for its 16,000 financial advisors," CNBC, March 14, 2023.

Production of tailored content at scale

Generative AI tools can draw on existing documents and data sets to substantially streamline content generation. These tools can create personalized marketing and sales content tailored to specific client profiles and histories as well as a multitude of alternatives for A/B testing. In addition, generative AI could automatically produce model documentation, identify missing documentation, and scan relevant regulatory updates to create alerts for relevant shifts.

Factors for banks to consider

When exploring how to integrate generative AI into operations, banks can be mindful of a number of factors:

- *The level of regulation for different processes.* These vary from unregulated processes such as customer service to heavily regulated processes such as credit risk scoring.
- *Type of end user.* End users vary widely in their expectations and familiarity with generative AI—for

example, employees compared with high-net-worth clients.

- *Intended level of work automation.* AI agents integrated through APIs could act nearly autonomously or as copilots, giving real-time suggestions to agents during customer interactions.
- *Data constraints.* While public data such as annual reports could be made widely available, there would need to be limits on identifiable details for customers and other internal data.

A generative AI bot trained on proprietary knowledge such as policies, research, and customer interaction could provide always-on, deep technical support.

Generative AI deployment could unlock potential value equal to 2.6 to 4.5 percent of annual revenues across the pharmaceutical and medical-product industries

Our analysis finds that generative AI could have a significant impact on the pharmaceutical and medical-product industries—from \$60 billion to \$110 billion annually. This big potential reflects the resource-intensive process of discovering new drug compounds. Pharma companies typically spend approximately 20 percent of revenues on R&D,¹ and the development of a new drug takes an average of ten to 15 years.

With this level of spending and timeline, improving the speed and quality of R&D can generate substantial value. For example, lead identification—a step in the drug discovery process in which researchers identify a molecule that would best address the target for a potential new drug—can take several months even with “traditional” deep learning techniques. Foundation models and generative AI can enable organizations to complete this step in a matter of weeks.

Generative AI use cases aligned to industry needs

Drug discovery involves narrowing the universe of possible compounds to those that could effectively treat specific conditions. Generative AI’s ability to process massive amounts of data and model options can accelerate output across several use cases:

Improve automation of preliminary screening

In the lead identification stage of drug development, scientists can use foundation models to automate the preliminary screening of chemicals in the search for those that will produce specific effects on drug targets. To start, thousands of cell cultures are tested and paired with images of the corresponding experi-

ment. Using an off-the-shelf foundation model, researchers can cluster similar images more precisely than they can with traditional models, enabling them to select the most promising chemicals for further analysis during lead optimization.

Enhance indication finding

An important phase of drug discovery involves the identification and prioritization of new indications—that is, diseases, symptoms, or circumstances that justify the use of a specific medication or other treatment, such as a test, procedure, or surgery. Possible indications for a given drug are based on a patient group’s clinical history and medical records, and they are then prioritized based on their similarities to established and evidence-backed indications.

Researchers start by mapping the patient cohort’s clinical events and medical histories—including potential diagnoses, prescribed medications, and performed procedures—from real-world data. Using foundation models, researchers can quantify clinical events, establish relationships, and measure the similarity between the patient cohort and evidence-backed indications. The result is a short list of indications that have a better probability of success in clinical trials because they can be more accurately matched to appropriate patient groups.

Pharma companies that have used this approach have reported high success rates in clinical trials for the top five indications recommended by a foundation model for a tested drug. This success has allowed these drugs to progress smoothly into Phase 3 trials, significantly accelerating the drug development process.

Additional factors to consider

Before integrating generative AI into operations, pharma executives should be aware of some factors that could limit their ability to capture its benefits:

- *The need for a human in the loop.* Companies may need to implement new quality checks on processes that shift from humans to generative AI, such as representative-generated emails, or more detailed quality checks on AI-assisted processes, such as drug discovery. The increasing need to verify whether generated content is based on fact or inference elevates the need for a new level of quality control.
- *Explainability.* A lack of transparency into the origins of generated content and traceability of root data could make it difficult to update models and scan them for potential risks; for instance, a generative AI solution for synthesizing scientific literature may not be able to point to the specific articles or quotes that led it to infer that a new treatment is very popular among physicians. The technology can also “hallucinate,” or generate responses that are obviously incorrect or inappropriate for the context. Systems need to be designed to point to specific articles or data sources, and then do human-in-the-loop checking.
- *Privacy considerations.* Generative AI’s use of clinical images and medical records could increase the risk that protected health information will leak, potentially violating regulations that require pharma companies to protect patient privacy.

¹ *Research and development in the pharmaceutical industry*, Congressional Budget Office, April 2021.



In this chapter, we have estimated the organizational value generative AI could deliver through use cases across industries and business functions, but the technology's potential is much greater. As it is embedded into tools used by every knowledge worker, its additional impact may be more diffuse but no less valuable than that associated with these use cases. Companies need to find ways to maximize the value created by the generative AI they deploy while also taking care to monitor and manage its impact on their workforces and society at large.



3

The generative AI future of work: Impacts on work activities, economic growth, and productivity

Technology has been changing the anatomy of work for decades. Over the years, machines have given human workers various “superpowers”; for instance, industrial-age machines enabled workers to accomplish physical tasks beyond the capabilities of their own bodies. More recently, computers have enabled knowledge workers to perform calculations that would have taken years to do manually.

These examples illustrate how technology can augment work through the automation of individual activities that workers would have otherwise had to do themselves. At a conceptual level, the application of generative AI may follow the same pattern in the modern workplace, although as we show later in this chapter, the types of activities that generative AI could affect, and the types of occupations with activities that could change, will likely be different as a result of this technology than for older technologies.

The McKinsey Global Institute began analyzing the impact of technological automation of work activities and modeling scenarios of adoption in 2017. At that time, we estimated that workers spent half of their time on activities that had the potential to be automated by

adapting technology that existed at that time, or what we call technical automation potential. We also modeled a range of potential scenarios for the pace at which these technologies could be adopted and affect work activities throughout the global economy.

Technology adoption at scale does not occur overnight. The potential of technological capabilities in a lab does not necessarily mean they can be immediately integrated into a solution that automates a specific work activity—developing such solutions takes time. Even when such a solution is developed, it might not be economically feasible to use if its costs exceed those of human labor. Additionally, even if economic incentives for deployment exist, it takes time for adoption to spread across the global economy. Hence, our adoption scenarios, which consider these factors together with the technical automation potential, provide a sense of the pace and scale at which workers' activities could shift over time.

Large-scale shifts in the mix of work activities and occupations are not unprecedented. Consider the work of a farmer today compared with what a farmer did just a few short years ago. Many farmers now access market information on mobile phones to determine when and where to sell their crops or download sophisticated modeling of weather patterns. From a more macro perspective, agricultural employment in China went from an 82 percent share of all workers in 1962 to 13 percent in 2013. Labor markets are also dynamic: millions of people leave their jobs every month in the United States.¹² But this does not minimize the challenges faced by individual workers whose lives are upended by these shifts, or the organizational or societal challenges of ensuring that workers have the skills to take on the work that will be in demand and that their incomes are sufficient to grow their standards of living.

Also, demographics have made such shifts in activities a necessity from a macroeconomic perspective. An economic growth gap has opened as a result of the slowing growth of the world's workforce. In some major countries, workforces have shrunk because populations are aging. Labor productivity will have to accelerate to achieve economic growth and enhance prosperity.

The analyses in this paper incorporate the potential impact of generative AI on today's work activities. The new capabilities of generative AI, combined with previous technologies and integrated into corporate operations around the world, could accelerate the potential for technical automation of individual activities and the adoption of technologies that augment the capabilities of the workforce. They could also have an impact on knowledge workers whose activities were not expected to shift as a result of these technologies until later in the future (see Box 3, "About the research").

**Labor productivity will
have to accelerate to
achieve economic growth
and enhance prosperity.**

Box 3

About the research

This analysis builds on the methodology we established in 2017. We began by examining the US Bureau of Labor Statistics O*Net breakdown of about 850 occupations into roughly 2,100 detailed work activities. For each of these activities, we scored the level of capability necessary to successfully perform the activity against a set of 18 capabilities that have the potential for automation (exhibit).

We also surveyed experts in the automation of each of these capabilities to estimate automation technologies' current performance level against each of these capabilities, as well as how the technology's performance might advance over time. Specifically, this year, we updated our assessments of technology's performance in cognitive, language, and social and emotional capabilities based on a survey of generative AI experts.

Based on these assessments of the technical automation potential of each detailed work activity at each point in time, we modeled potential scenarios for the adoption of work automation around the world. First, we estimated a range of time to implement a solution that could automate each specific detailed work activity, once all the capability requirements were met by the state of technology development. Second, we estimated a range of potential costs for this technology when it is first introduced, and then declining over time, based on historical precedents. We modeled the beginning of adoption for a specific detailed work activity in a particular occupation in a country (for 47 countries, accounting for more than 80 percent of the global workforce) when the cost of the automation technology reaches parity with the cost of human labor in that occupation.

Based on a historical analysis of various technologies, we modeled a range of adoption timelines from eight to 27 years between the beginning of adoption and its plateau, using sigmoidal curves (S-curves). This range implicitly accounts for the many factors that could affect the pace at which adoption occurs, including regulation, levels of investment, and management decision making within firms.

The modeled scenarios create a time range for the potential pace of automating current work activities. The "earliest" scenario flexes all parameters to the extremes of plausible assumptions, resulting in faster automation development and adoption, and the "latest" scenario flexes all parameters in the opposite direction. The reality is likely to fall somewhere between the two.

Exhibit

Our analysis assesses the potential for technical automation across some 2,100 activities and 18 capabilities.

~850 occupations	~2,100 activities assessed across all occupations	Capability requirements	
Examples			
 Retail salespeople  Health practitioners  Food and beverage service workers  Teachers	Example: Retail activities • Answer questions about products and services • Greet customers • Clean and maintain work areas • Demonstrate product features • Process sales and transactions	Sensory ✔ Sensory perception Cognitive ✔ Retrieving information ✔ Recognizing known patterns and categories (supervised learning) ✔ Generating novel patterns and categories ✔ Logical reasoning and problem solving ✔ Optimizing and planning ✔ Creativity ✔ Articulating/display output ✔ Coordination with multiple agents	Physical ✔ Fine motor skills and dexterity ✔ Gross motor skills ✔ Navigation ✔ Mobility Natural-language processing ✔ Understanding natural language ✔ Generating natural language Social ✔ Social and emotional sensing ✔ Social and emotional reasoning ✔ Social and emotional output

Source: McKinsey Global Institute analysis

Accelerating the technical potential to transform knowledge work

Based on developments in generative AI, technology performance is now expected to match median human performance and reach top quartile human performance earlier than previously estimated across a wide range of capabilities (Exhibit 6). For example, MGI previously identified 2027 as the earliest year when median human performance for natural-language understanding might be achieved in technology, but in this new analysis, the corresponding point is 2023.

Exhibit 6

As a result of generative AI, experts assess that technology could achieve human-level performance in some technical capabilities sooner than previously thought.

Technical capabilities, level of human performance achievable by technology



¹Comparison made on the business-related tasks required from human workers. Please refer to technical appendix for detailed view of performance rating methodology.

Source: McKinsey Global Institute occupation database; McKinsey analysis

McKinsey & Company


As a result of these reassessments of technology capabilities due to generative AI, the total percentage of hours that could theoretically be automated by integrating technologies that exist today has increased from about 50 percent to 60–70 percent. The technical potential curve is quite steep because of the acceleration in generative AI's natural-language capabilities (Exhibit 7).

Interestingly, the range of times between the early and late scenarios has compressed compared with the expert assessments in 2017, reflecting a greater confidence that higher levels of technological capabilities will arrive by certain time periods.

Exhibit 7

The advent of generative AI has pulled forward the potential for technical automation.

Technical automation potentials by scenario, %



¹Includes data from 47 countries, representing about 80% of employment across the world. 2017 estimates are based on the activity and occupation mix from 2016. Scenarios including generative AI are based on the 2021 activity and occupation mix.

²Early and late scenarios reflect the ranges provided by experts (see Exhibit 6).

Source: McKinsey Global Institute analysis

McKinsey & Company

Adoption lags behind technical automation potential

Our analysis of adoption scenarios accounts for the time required to integrate technological capabilities into solutions that can automate individual work activities; the cost of these technologies compared with that of human labor in different occupations and countries around the world; and the time it has taken for technologies to diffuse across the economy. With the acceleration in technical automation potential that generative AI enables, our scenarios for automation adoption have correspondingly accelerated. These scenarios encompass a wide range of outcomes, given that the pace at which solutions will be developed and adopted will vary based on decisions that will be made on investments,

deployment, and regulation, among other factors. But they give an indication of the degree to which the activities that workers do each day may shift.


As an example of how this might play out in a specific occupation, consider postsecondary English language and literature teachers, whose detailed work activities include preparing tests and evaluating student work. With generative AI's enhanced natural-language capabilities, more of these activities could be done by machines, perhaps initially to create a first draft that is edited by teachers but perhaps eventually with far less human editing required. This could free up time for these teachers to spend more time on other work activities, such as guiding class discussions or tutoring students who need extra assistance.

Our previously modeled adoption scenarios suggested that 50 percent of time spent on 2016 work activities would be automated sometime between 2035 and 2070, with a midpoint scenario around 2053. Our updated adoption scenarios, which account for developments in generative AI, models the time spent on 2023 work activities reaching 50 percent automation between 2030 and 2060, with a midpoint of 2045—an acceleration of roughly a decade compared with the previous estimate (Exhibit 8).¹³

Exhibit 8

The midpoint scenario at which automation adoption could reach 50 percent of time spent on current work activities has accelerated by a decade.

Global automation of time spent on current work activities,¹ %



¹Includes data from 47 countries, representing about 80% of employment across the world. 2017 estimates are based on the activity and occupation mix from 2016. Scenarios including generative AI are based on the 2021 activity and occupation mix.

²Early scenario: aggressive scenario for all key model parameters (technical automation potential, integration timelines, economic feasibility, and technology diffusion rates).

³Late scenario: parameters are set for later adoption potential.
Source: McKinsey Global Institute analysis

Different countries, different pace of adoption

Adoption is also likely to be faster in developed countries, where wages are higher and thus the economic feasibility of adopting automation occurs earlier. Even if the potential for technology to automate a particular work activity is high, the costs required to do so have to be compared with the cost of human wages. In countries such as China, India, and Mexico, where wage rates are lower, automation adoption is modeled to arrive more slowly than in higher-wage countries (Exhibit 9).


Exhibit 9

Automation adoption is likely to be faster in developed economies, where higher wages will make it economically feasible sooner.


Automation adoption by scenario for select countries, %

United States Germany Japan France China Mexico India

Early scenario¹



Late scenario²



¹Early scenario: aggressive scenario for all key model parameters (technical automation potential, integration timelines, economic feasibility, and technology diffusion rates.).

²Late scenario: parameters are set for the later adoption potential.

Source: McKinsey Global Institute analysis

McKinsey & Company

Our analyses of generative AI's impact on work activities and the pace of automation adoption rely on several assumptions and sensitivities (see Box 4, "Limitations of our analyses, key assumptions, and sensitivities").

Box 4

Limitations of our analyses, key assumptions, and sensitivities

This analysis considers the potential for automation only of current work activities and occupations. It does not account for how those work activities may shift over time or forecast new activities and occupations.¹ Also, the analysis accounts solely for first-order effects. It does not take into account how labor rates could change, and it does not model changes in labor force participation rates or other general equilibrium effects. That said, while these models account for the time it may take for technology to be adopted across an economy, technologies could be adopted much more rapidly in an individual organization. Other research may reach different conclusions.

Our assessments of technology capabilities are based on the best estimates of experts involved in

developing automation technologies. These assessments could change over time, as they have changed since 2017.

The technology adoption curves are based on historical findings that technologies take eight to 27 years from commercial availability to reach a plateau in adoption. Some argue that the adoption of generative AI will be faster due to the ease of deploying these technologies. That said, the case for a minimum of eight years in our earliest scenario for reaching a global plateau in adoption accounts for the pace of adoption of other technologies that have arguably had a faster adoption potential—for example, social networking as a consumer technology that faced no barriers in enterprise change management. Our scenario also accounts for the significant

role of small and midsize enterprises around the world, in addition to the challenges of incorporating and managing change in larger organizations.

In addition, this analysis does not assume that the scale of work automation equates directly to job losses. Like other technologies, generative AI typically enables individual activities within occupations to be automated, not entire occupations. Historically, the activities in many occupations have shifted over time as certain activities are automated. However, organizations may decide to realize the benefits of increased productivity by reducing employment in some job categories, a possibility we cannot rule out.

¹ David Autor et al., *New frontiers: The origins and content of new work, 1940–2018*, National Bureau of Economic Research working paper number 30389, August 2022; Jeffrey Lin, “Technological adaptation, cities, and new work,” *Review of Economics and Statistics*, May 2011, Volume 93, Number 2.

Generative AI is likely to have the biggest impact on knowledge work, particularly activities involving decision making and collaboration, which previously had the lowest potential for automation.

Generative AI's potential impact on knowledge work

Previous generations of automation technology were particularly effective at automating data management tasks related to collecting and processing data. Generative AI's natural-language capabilities increase the automation potential of these types of activities somewhat. But its impact on more physical work activities shifted much less, which isn't surprising because its capabilities are fundamentally engineered to do cognitive tasks.


As a result, generative AI is likely to have the biggest impact on knowledge work, particularly activities involving decision making and collaboration, which previously had the lowest potential for automation (Exhibit 10). Our estimate of the technical potential to automate the application of expertise jumped 34 percentage points, while the potential to automate management and develop talent increased from 16 percent in 2017 to 49 percent in 2023.

Generative AI's ability to understand and use natural language for a variety of activities and tasks largely explains why automation potential has risen so steeply. Some 40 percent of the activities that workers perform in the economy require at least a median level of human understanding of natural language.

Exhibit 10

Generative AI could have the biggest impact on collaboration and the application of expertise, activities that previously had a lower potential for automation.

Overall technical automation potential, comparison in midpoint scenarios, % in 2023



Note: Figures may not sum, because of rounding.

¹Previous assessment of work automation before the rise of generative AI.

²Applying expertise to decision making, planning, and creative tasks.

³Managing and developing people.

⁴Performing physical activities and operating machinery in unpredictable environments.

⁵Performing physical activities and operating machinery in predictable environments.

Source: McKinsey Global Institute analysis


McKinsey & Company

As a result, many of the work activities that involve communication, supervision, documentation, and interacting with people in general have the potential to be automated by generative AI, accelerating the transformation of work in occupations such as education and technology, for which automation potential was previously expected to emerge later (Exhibit 11).

Exhibit 11

Advances in technical capabilities could have the most impact on activities performed by educators, professionals, and creatives.

Impact of generative AI on technical automation potential in midpoint scenario, 2023



Note: Figures may not sum, because of rounding.

¹Previous assessment of work automation before the rise of generative AI.

²Includes data from 47 countries, representing about 80% of employment across the world.

Source: McKinsey Global Institute analysis


McKinsey & Company

Labor economists have often noted that the deployment of automation technologies tends to have the most impact on workers with the lowest skill levels, as measured by educational attainment, or what is called skill biased. We find that generative AI has the opposite pattern—it is likely to have the most incremental impact through automating some of the activities of more-educated workers (Exhibit 12).

Exhibit 12

Generative AI increases the potential for technical automation most in occupations requiring higher levels of educational attainment.

Impact of generative AI on technical automation potential in midpoint scenario, 2023



¹Previous assessment of work automation before the rise of generative AI.
Source: McKinsey Global Institute analysis

McKinsey & Company

Another way to interpret this result is that generative AI will challenge the attainment of multiyear degree credentials as an indicator of skills, and others have advocated for taking a more skills-based approach to workforce development in order to create more equitable, efficient workforce training and matching systems.¹⁴ Generative AI could still be described as skill-biased technological change, but with a different, perhaps more granular, description of skills that are more likely to be replaced than complemented by the activities that machines can do.

Previous generations of automation technology often had the most impact on occupations with wages falling in the middle of the income distribution. For lower-wage occupations, making a case for work automation is more difficult because the potential benefits of automation compete against a lower cost of human labor. Additionally, some of the tasks performed in lower-wage occupations are technically difficult to automate—for example, manipulating fabric or picking delicate fruits. Some labor economists have observed a


“hollowing out of the middle,” and our previous models have suggested that work automation would likely have the biggest midterm impact on lower-middle-income quintiles.

However, generative AI's impact is likely to most transform the work of higher-wage knowledge workers because of advances in the technical automation potential of their activities, which were previously considered to be relatively immune from automation (Exhibit 13).

Exhibit 13

Generative AI could have the biggest impact on activities in high-wage jobs; previously, automation's impact was highest in lower-middle-income quintiles.

Automation adoption per wage quintile, % in 2030, midpoint scenario



¹Previous assessment of work automation before the rise of generative AI.
Source: McKinsey Global Institute analysis

Generative AI could propel higher productivity growth


Global economic growth was slower from 2012 to 2022 than in the two preceding decades.¹⁵ Although the COVID-19 pandemic was a significant factor, long-term structural challenges—including declining birth rates and aging populations—are ongoing obstacles to growth.

Declining employment is among those obstacles. Compound annual growth in the total number of workers worldwide slowed from 2.5 percent in 1972–82 to just 0.8 percent in 2012–22, largely because of aging. In many large countries, the size of the workforce is already declining.¹⁶ Productivity, which measures output relative to input, or the value of goods and services produced divided by the amount of labor, capital, and other resources required to produce them, was the main engine of economic growth in the three decades from 1992 to 2022 (Exhibit 14). However, since then, productivity growth has slowed in tandem with slowing employment growth, confounding economists and policy makers.¹⁷

Exhibit 14

Productivity growth, the main engine of GDP growth over the past 30 years, slowed down in the past decade.

Real GDP growth contribution of employment and productivity growth, 1972–2022,
global GDP growth, CAGR, %



Source: Conference Board Total Economy database; McKinsey Global Institute analysis

McKinsey & Company


The deployment of generative AI and other technologies could help accelerate productivity growth, partially compensating for declining employment growth and enabling overall economic growth. Based on our estimates, the automation of individual work activities enabled by these technologies could provide the global economy with an annual productivity boost of 0.2 to 3.3 percent from 2023 to 2040 depending on the rate of automation adoption—with generative AI contributing to 0.1 to 0.6 percentage points of that growth—but only if individuals affected by the technology were to shift to other work activities that at least match their 2022 productivity levels (Exhibit 15). In some cases, workers will stay in the same occupations, but their mix of activities will shift; in others, workers will need to shift occupations.

Exhibit 15

Generative AI could contribute to productivity growth if labor hours can be redeployed effectively.

Productivity impact from automation by scenario, 2022–40, CAGR,¹ %

■ Without generative AI² ■ Additional with generative AI



Note: Figures may not sum, because of rounding.

¹Based on the assumption that automated work hours are reintegrated in work at productivity level of today.


²Previous assessment of work automation before the rise of generative AI.

³Based on 47 countries, representing about 80% of world employment.

Source: Conference Board Total Economy database; Oxford Economics; McKinsey Global Institute analysis

McKinsey & Company

The capabilities of generative AI vastly expand the pool of work activities with the potential for technical automation. That in turn has sped up the pace at which automation may be deployed and expanded the types of workers who will experience its impact. Like other technologies, its ability to take on routine tasks and work can increase human productivity, which has grown at a below-average rate for almost 20 years.¹⁸ It can also offset the impact of aging, which is beginning to put a dent in workforce growth for many of the world's major economies. But to achieve these benefits, a significant number of workers will need to substantially change the work they do, either in their existing occupations or in new ones. They will also need support in making transitions to new activities.





4

Considerations for businesses and society

History has shown that new technologies have the potential to reshape societies. Artificial intelligence has already changed the way we live and work—for example, it can help our phones (mostly) understand what we say, or draft emails. Mostly, however, AI has remained behind the scenes, optimizing business processes or making recommendations about the next product to buy. The rapid development of generative AI is likely to significantly augment the impact of AI overall, generating trillions of dollars of additional value each year and transforming the nature of work.

But the technology could also deliver new and significant challenges. Stakeholders must act—and quickly, given the pace at which generative AI could be adopted—to prepare to address both the opportunities and the risks. Risks have already surfaced, including concerns about the content that generative AI systems produce: Will they infringe upon intellectual property due to “plagiarism” in the training data used to create foundation models? Will the answers that LLMs produce when questioned be accurate, and can they be explained? Will the content generative AI creates be fair or biased in ways that users do not want by, say, producing content that reflects harmful stereotypes?

There are economic challenges too: the scale and the scope of the workforce transitions described in this report are considerable. In the midpoint adoption scenario, about a quarter to a third of work activities could change in the coming decade. The task before us is to manage the potential positives and negatives of the technology simultaneously (for more about the potential risks of generative AI, see Box 5). Here are some of the critical questions we will need to address while balancing our enthusiasm for the potential benefits of the technology with the new challenges it can introduce.

Box 5

Using generative AI responsibly

Generative AI poses a variety of risks.

Stakeholders will want to address these risks from the start.

Fairness: Models may generate algorithmic bias due to imperfect training data or decisions made by the engineers developing the models.

Intellectual property (IP): Training data and model outputs can generate significant IP risks, including infringing on copyrighted, trademarked, patented, or otherwise legally protected materials. Even when using a provider's generative AI tool, organizations will need to understand what data went into training and how it's used in tool outputs.

Privacy: Privacy concerns could arise if users input information that later ends up in model outputs in a form that makes

individuals identifiable. Generative AI could also be used to create and disseminate malicious content such as disinformation, deepfakes, and hate speech.

Security: Generative AI may be used by bad actors to accelerate the sophistication and speed of cyberattacks. It also can be manipulated to provide malicious outputs. For example, through a technique called prompt injection, a third party gives a model new instructions that trick the model into delivering an output unintended by the model producer and end user.

Explainability: Generative AI relies on neural networks with billions of parameters, challenging our ability

to explain how any given answer is produced.

Reliability: Models can produce different answers to the same prompts, impeding the user's ability to assess the accuracy and reliability of outputs.

Organizational impact: Generative AI may significantly affect the workforce, and the impact on specific groups and local communities could be disproportionately negative.

Social and environmental impact: The development and training of foundation models may lead to detrimental social and environmental consequences, including an increase in carbon emissions (for example, training one large language model can emit about 315 tons of carbon dioxide).¹

¹ Ananya Ganesh, Andrew McCallum, and Emma Strubell, "Energy and policy considerations for deep learning in NLP," *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, June 5, 2019.

Companies and business leaders

How can companies move quickly to capture the potential value at stake highlighted in this report, while managing the risks that generative AI presents?

How will the mix of occupations and skills needed across a company's workforce be transformed by generative AI and other artificial intelligence over the coming years? How will a company enable these transitions in its hiring plans, retraining programs, and other aspects of human resources?

Do companies have a role to play in ensuring the technology is not deployed in "negative use cases" that could harm society?

How can businesses transparently share their experiences with scaling the use of generative AI within and across industries—and also with governments and society?

Policy makers

What will the future of work look like at the level of an economy in terms of occupations and skills? What does this mean for workforce planning?

How can workers be supported as their activities shift over time? What retraining programs can be put in place? What incentives are needed to support private companies as they invest in human capital? Are there earn-while-you-learn programs such as apprenticeships that could enable people to retrain while continuing to support themselves and their families?

What steps can policy makers take to prevent generative AI from being used in ways that harm society or vulnerable populations?

Can new policies be developed and existing policies amended to ensure human-centric AI development and deployment that includes human oversight and diverse perspectives and accounts for societal values?

Individuals as workers, consumers, and citizens

How concerned should individuals be about the advent of generative AI? While companies can assess how the technology will affect their bottom lines, where can citizens turn for accurate, unbiased information about how it will affect their lives and livelihoods?

How can individuals as workers and consumers balance the conveniences generative AI delivers with its impact in their workplaces?

Can citizens have a voice in the decisions that will shape the deployment and integration of generative AI into the fabric of their lives?

Technological innovation can inspire equal parts awe and concern. When that innovation seems to materialize fully formed and becomes widespread seemingly overnight, both responses can be amplified. The arrival of generative AI in the fall of 2022 was the most recent example of this phenomenon, due to its unexpectedly rapid adoption as well as the ensuing scramble among companies and consumers to deploy, integrate, and play with it.

All of us are at the beginning of a journey to understand this technology's power, reach, and capabilities. If the past eight months are any guide, the next several years will take us on a roller-coaster ride featuring fast-paced innovation and technological breakthroughs that force us to recalibrate our understanding of AI's impact on our work and our lives. It is important to properly understand this phenomenon and anticipate its impact. Given the speed of generative AI's deployment so far, the need to accelerate digital transformation and reskill labor forces is great.

These tools have the potential to create enormous value for the global economy at a time when it is pondering the huge costs of adapting and mitigating climate change. At the same time, they also have the potential to be more destabilizing than previous generations of artificial intelligence. They are capable of that most human of abilities, language, which is a fundamental requirement of most work activities linked to expertise and knowledge as well as a skill that can be used to hurt feelings, create misunderstandings, obscure truth, and incite violence and even wars.

We hope this research has contributed to a better understanding of generative AI's capacity to add value to company operations and fuel economic growth and prosperity as well as its potential to dramatically transform how we work and our purpose in society. Companies, policy makers, consumers, and citizens can work together to ensure that generative AI delivers on its promise to create significant value while limiting its potential to upset lives and livelihoods. The time to act is now.¹⁹

Endnotes

- 1 “A future that works: Automation, employment, and productivity,” McKinsey Global Institute, January 12, 2017.
- 2 Ryan Morrison, “Compute power is becoming a bottleneck for developing AI. Here’s how you clear it.,” *Tech Monitor*, updated March 17, 2023.
- 3 “Introducing ChatGPT,” OpenAI, November 30, 2022; “GPT-4 is OpenAI’s most advanced system, producing safer and more useful responses,” OpenAI, accessed June 1, 2023.
- 4 “Introducing Claude,” Anthropic PBC, March 14, 2023; “Introducing 100K Context Windows,” Anthropic PBC, May 11, 2023.
- 5 Emma Roth, “The nine biggest announcements from Google I/O 2023,” *The Verge*, May 10, 2023.
- 6 Pitchbook.
- 7 Ibid.
- 8 Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, *Generative AI at work*, National Bureau of Economic Research working paper number 31161, April 2023.
- 9 Peter Cihon et al., *The impact of AI on developer productivity: Evidence from GitHub Copilot*, Cornell University arXiv software engineering working paper, arXiv:2302.06590, February 13, 2023.
- 10 Michael Nuñez, “Google and Replit join forces to challenge Microsoft in coding tools,” *VentureBeat*, March 28, 2023.
- 11 Joe Coscarelli, “An A.I. hit of fake ‘Drake’ and ‘The Weeknd’ rattles the music world,” *New York Times*, updated April 24, 2023.
- 12 “Job openings and labor turnover survey,” US Bureau of Labor Statistics, accessed June 6, 2023.
- 13 The comparison is not exact because the composition of work activities between 2016 and 2023 has changed; for example, some automation has occurred during that time period.
- 14 A more skills-based approach to workforce development predates the emergence of generative AI.
- 15 *Global economic prospects*, World Bank, January 2023.
- 16 Yaron Shamir, “Three factors contributing to fewer people in the workforce,” *Forbes*, April 7, 2022.
- 17 “The U.S. productivity slowdown: an economy-wide and industry-level analysis,” *Monthly Labor Review*, US Bureau of Labor Statistics, April 2021; Kweilin Ellingrud, “Turning around the productivity slowdown,” McKinsey Global Institute, September 13, 2022.
- 18 “Rekindling US productivity for a new era,” McKinsey Global Institute, February 16, 2023.
- 19 The research, analysis, and writing in this report was entirely done by humans.



Appendix

I. Scope of the investment landscape

For data on investment flows into generative AI, we relied on Pitchbook.

For data on investments in artificial intelligence overall, we referred to *The AI Index 2023 annual report* by Stanford University's Institute for Human-Centered AI.¹

II. How we sized the use case value potential

Overall objective

The objective was to estimate the potential economic impact of generative AI and foundation models using a bottom-up assessment of the most relevant use cases across business functions and industries. The resulting analysis approximates the potential value of generative AI and foundation models in terms of productivity, or the equivalent amount by which the technologies could reduce the global functional spending required to maintain current revenue levels.

Summary of our approach

Our micro-to-macro use-case-based approach to estimating the potential impact of generative AI included the following:

How we estimated the impact of generative AI across industries. We identified and cataloged generative AI and foundation model use cases with input from experts from McKinsey's industry and functional practices. While the resulting generative AI and foundation model use case database is not necessarily exhaustive, it was meant to be as comprehensive as possible.

Business function overview

Business function	Definition
Customer operations	Activities related to customer care, such as call centers
Software development	Activities related to designing, coding, testing, and maintaining software programs, applications, and systems, such as enterprise resource planning tools and other internal tools, that meet specific business or customer needs
Product R&D	Activities related to new products or services, such as market and academic research, ideation, and simulations used in the early stages of product development, and heavy simulations, prototyping, and testing used later in the development process
Legal, risk, and compliance	Activities related to risk and compliance, such as labor relations, litigation support, and contract creation

¹ *The AI Index 2023 annual report*, Institute for Human-Centered AI, Stanford University, April 2023.

Marketing and sales	Activities related to B2B and B2C marketing, for example, market research, creative, marketing strategy, and sales, customer preparation, and interaction support, including pricing analytical support, price scraping, and item matching
IT	Activities related to internal information tech systems, such as administrative and IT help desk, but excluding software engineering for internal solutions (captured within software engineering)
Talent and organization (including HR)	Activities related to organizational performance and talent management; for example, organizational health assessment, recruiting, learning and development, and human resources
Finance and strategy	Activities related to financial operations, such as general accounting, financial planning and analysis, and accounts payable and receivable, as well as internal strategy functions such as market intelligence and strategic planning

Estimate generative AI's impact in individual use cases. We collected expert inputs, the results of McKinsey internal experiments, and published research to estimate the potential quantitative range of impact (in both cost savings and revenue uplift) and to gain qualitative findings from functions within individual use cases.² Where generative AI and foundation models are assumed to increase revenue, we recast the effect as an increase in productivity that would be the equivalent of the reduced level of spending required to maintain the same level of output (thus enabling comparability with cost reductions).

Our analysis of use cases only examines the direct impact of generative AI on productivity. They do not incorporate secondary benefits, such as the economic impact of hiring a more capable employee or hiring employees more rapidly.

Estimate the impact across industries

To estimate the impact of generative AI and foundation models across industries, we scaled our analysis from a functional lens to an industry lens for each industry, assessing the weight of functional costs. For instance, customer operations costs are higher in the telecom industry than in aerospace.

For each use case, the relevant costs were defined to account only for activities for which generative AI or foundation models were likely to deliver productivity gains. For example, the cost base against which we assessed productivity gains for the marketing function excludes ad buying costs because the use cases we assessed for generative AI in that function would have no productivity impact on the “ad buying” activity.

- A. Global revenues for each industry in 2022 were sourced from IHS Markit and Oxford Economics.
- B. We estimated the functional cost for a function as a percentage of total revenue in an industry based on published data, industry experts, and McKinsey benchmarks.

² *Generative AI at work*, April 2023.

- C. For each function, we estimated the relevant costs (considering only underlying activities affected by our use cases) as a portion of overall spending on a function, again informed by published data, industry experts, and McKinsey benchmarks.
- D. For each use case, we quantified the potential impact of generative AI and foundation models based on the impact (either cost or revenue) as a function of the relevant addressable spending.
- E. We then calculated the impact by industry by aggregating the technology's impact on use cases in an industry across functions.

Limitations of our impact analysis

For each use case, our analysis of the impact of generative AI and foundation models draws a conservative base case and a more accelerated upside potential. The estimates are based on existing data and experiences, and these estimates could be updated over time. Additionally, the list of use cases is not exhaustive but is as comprehensive as possible, given our methods.

The estimates in this report should be treated as directional rather than precise, given the nature of the technology and the wide range of uncertainty involved in the future development of generative AI. We welcome challenges to our analysis as well as additional inputs that would refine and enhance it.

Comparison with the 2018 report *Notes from the AI frontier*:

Applications and value of deep learning

To estimate the incremental impact that generative AI and foundation models could have in comparison to the value of artificial intelligence overall, we updated our 2018 estimates of the value that could be created through deployment of advanced analytics and previous generations of artificial intelligence.³

This research incorporates 2022 updates to baseline financial variables such as industry revenues, enabling updates to the total potential economic impact of advanced analytics and AI exclusive of generative AI and foundation models.

III. How we estimated the impact of generative AI on the potential for technical automation

This report continues and adapts the methodology and findings of the January 2017 McKinsey Global Institute report, *A future that works: Automation, employment, and productivity*.⁴ We recommend that readers refer to the technical appendix of that report for our full methodology. Here, we outline updates we have incorporated in the automation adoption estimation process and a brief overview of the steps involved in assessing automation adoption, together with methodology to estimate the impact of automation on work hours (represented as full-time equivalents, or FTEs, where one FTE equals 2,080 hours) and GDP impact. (For additional research on the impact of automation on work transformation, please see Box A1, "Overview of select recent studies on the impact of generative AI on work automation").

³ "Notes from the AI frontier," April 17, 2018.

⁴ "Harnessing automation for a future that works," January 12, 2017.

Overview of select recent studies on the impact of generative AI on work automation

Goldman Sachs

Citation	Joseph Briggs et al., <i>The potentially large effects of artificial intelligence on economic growth</i> , Goldman Sachs, March 26 2023
Unit of analysis	Occupations; industries
Scope	United States and Europe (with extrapolation globally)
Approach summary	- Estimate the share of total work exposed to labor-saving automation by AI by occupation and industry. - Thirteen work activities are classified as exposed to AI automation based on a review on the probable use cases of Generative AI, and a share of each occupation's total workload that AI has the potential to replace is estimated by applying the O*Net "level" scale and taking an importance- and complexity-weighted average of essential work tasks.
Data collection	O*NET and European ESCO databases were used to obtain information about task contents and occupations.
Key relevant findings	- In the United States and Europe, two-thirds of jobs are being exposed to some degree of AI automation; 25 percent of current work tasks could be automated. Extrapolating globally, 18 percent of work could be automated by AI, with scenarios ranging from 15 to 35 percent, depending on different levels of AI capabilities. - Seven percent of US employment would be substituted by AI, and 300 million jobs globally would be exposed to automation, assuming jobs for which at least 50 percent of importance- and complexity-weighted tasks are exposed to automation are likely to be substituted by AI, 10 to 49 percent are likely to be complemented, and less than 10 percent are unlikely to be affected. - Generative AI could raise annual US labor productivity growth by 1.5 percentage points over a ten-year period, assuming that 7 percent of workers are fully displaced but able to find employment at slightly less productive positions, that partially exposed workers increase their productivity, and that roughly half of firms adopt generative AI during that period. Scenarios of productivity growth range from 0.3 to 3.0 percentage points depending on the difficulty level of tasks generative AI could perform, how many jobs are automated, and the speed of adoption. - Extrapolating globally, annual productivity growth from generative AI could be 1.4 percentage points over ten years.

Open AI, OpenResearch, University of Pennsylvania

Citation	Tyna Eloundou et al., <i>GPTs are GPTs: An early look at the labor market impact potential of large language models</i> , Cornell University arXiv economics working paper, arXiv:2303.10130, March 2023
Unit of analysis	Skills; occupation groups; industries
Scope	United States
Approach summary	- Investigate implications of large language models (LLMs) on the United States labor market using three exposure categories (exposure being a measure of whether access to an LLM would reduce the time required to complete a specific detailed work activity or task by at least 50 percent): 1 is no exposure, 2 is direct exposure, and 3 is LLM+ exposure (that is, additional software could be developed on top of the LLM) to reach 50 percent. - Exposure was assessed by human annotators and GPT-4.
Data collection	- The working paper used the O*NET database (1,016 occupations, 2,087 detailed work activities [DWA], and 19,265 tasks). DWAs were first aggregated to task level then occupation level. - Employment and wage data were obtained from the 2020 and 2021 occupational employment series provided by the US Bureau of Labor Statistics.
Key relevant findings	- About 80 percent of the US workforce could have at least 10 percent of their work tasks exposed to LLMs (the time required to complete these tasks reduced by at least 50 percent). - Nineteen percent of workers may see at least 50 percent of their tasks exposed to LLMs. - About 15 percent of all worker tasks in the United States could be performed significantly faster at the same level of quality with LLMs, but this share increases to 47 to 56 percent when incorporating software and tooling built on top of LLMs. - Effects of LLMs are relevant to all wage levels, but high-income occupations and occupations requiring at least a bachelor's degree may be more exposed than others.

Box A1 (continued)

Princeton University, University of Pennsylvania, New York University

Citation	Edward W. Felten, Manav Raj, and Robert Seamans, <i>Occupational heterogeneity in exposure to generative AI</i> , SSRN, April 2023
Unit of analysis	Occupations
Scope	United States
Approach summary	- Estimate exposure scores of the occupations most exposed to advances in AI language modeling and AI image generation in US workforce. - Ten AI applications (such as image generation, language modeling, and real-time video games) are linked to 52 human abilities (such as oral expression) through crowdsourced assessments of relatedness, which are mapped to approximately 800 occupations to create exposure scores for language modeling and image generation.
Data collection	- O*NET database provides mapping from occupations to human abilities, including prevalence and importance scores. - American Community Survey five-year estimates from 2021 provided by IPUMS were used to provide data on demographics.
Key relevant findings	- The average occupational exposure to language modeling is higher than to image generation. - Occupations most exposed to language modeling require communication and language-based abilities, while those most exposed to image generation require visual or spatial abilities. - Highly-educated, highly-paid, white-collar occupations may be most exposed to generative AI. - Strong, positive correlation between generative AI exposure and median salaries, required education levels, and the presence of creative abilities within an occupation. - There is positive correlation between generative AI exposure and percent of female, White, and Asian representation in occupations in US workforce. There is negative correlation between generative AI exposure and percent of male, Black, and Hispanic representation.

McKinsey & Company

Citation	<i>The economic potential of generative AI: The next productivity frontier</i> , McKinsey Global Institute, June 2023
Unit of analysis	Function; industries; occupations
Scope	Forty-seven countries covering 80 percent of global workforce
Approach summary	- Estimate when technologies including generative AI will reach median and top quartile human performance against 18 different capabilities through expert assessments. - Map these assessments to DWAs and occupations to develop scenarios of technical automation potential increasing over time - Model scenarios for automation adoption for DWAs, occupations and countries that include technical automation potential, solution integration timelines, economic feasibility versus wage rates, and technology diffusion rates.
Data collection	- Occupation and detailed work activity data was sourced from O*NET. - Employment and wage data was sourced from national statistical agencies (such as the US Bureau of Labor Statistics).
Key relevant findings	- The capabilities of generative AI have increased the share of time spent on work activities that theoretically could be automated by adapting technologies available in 2023 from about half to two-thirds of all working time. - The pace of adoption is likely to accelerate, with estimates that half of today's work activities could be automated in scenarios that range from 2030 to 2060, with a midpoint in 2045—or roughly a decade earlier than the midpoint estimate produced by our 2017 scenarios. - Experts' assessment shows much of this acceleration in the potential for technical automation is due to generative AI's increased natural language capabilities. - Unlike most automation in the past, the impact of generative AI will fall most heavily on occupations requiring higher levels of education and commanding higher wages. - Generative AI could increase labor productivity by 0.1 percent to 0.6 percent annually over the next ten to 20 years. When combined with other technologies, automation overall could contribute 0.2 to 3.3 percent annually to productivity growth, assuming labor is redeployed at today's productivity levels and not including general equilibrium effects.

Source: US Bureau of Labor Statistics O*NET; McKinsey analysis

McKinsey & Company

We assessed the technical potential for automation across the global economy through an analysis of the component activities of each occupation. Our analysis covers 47 countries representing more than 80 percent of the global economy. We used the US Bureau of Labor Statistics O*Net data set that maps about 850 occupations to approximately 2,100 detailed work activities.

Each detailed work activity was assessed against 18 capabilities that could potentially be automated based on the level of performance necessary to successfully perform that activity. These assessments were informed by academic research, McKinsey expertise, and industry experts. We defined four possible levels of requirement for each capability, ranging from not required to top-quartile human performance. Please refer to exhibits A1 to A4 of the technical appendix in the January 2017 McKinsey Global Institute report, *A future that works: Automation, employment, and productivity* for more details on the capabilities and the four requirement levels.⁵

Assigning the required level of capabilities to activities

We used a machine learning algorithm to score the approximately 2,100 work activities in relation to the 18 performance capabilities. To train the algorithm, we devised a list of keywords with input from experts. The algorithm scores each activity by matching keywords from the capability to the activity title. These assessments were checked manually, and where we found anomalies, special requirements, or a need for nuance, we made adjustments—for example, in assessing the level of capabilities needed to navigate in extreme weather or on uneven surfaces and other unpredictable settings, or the different physical capabilities required by a kindergarten teacher compared with a middle school teacher.

1. Modeling of automation adoption timelines

Our adoption model assesses the automation development and adoption timeline at activity level for more than 850 occupations across 19 sectors and 47 countries that represent more than 80 percent of the global economy. We divide the adoption process into four phases: technical feasibility, solution development, economic feasibility, and end-user adoption.

Technical feasibility

For a detailed work activity to be automated, technology must exist that performs at the required level for each of the 18 capabilities. To update the progression scenarios for the performance capabilities over time, we surveyed experts in generative AI. We also conducted interviews with industry leaders and academic experts. Based on these assessments, we updated the expected time frames to reach each level of performance for each capability. In our analysis and interaction with experts, we focused on the following capabilities that could be affected by generative AI:

- Natural-language understanding
- Natural-language generation
- Social and emotional reasoning
- Emotional and social output
- Social and emotional sensing
- Coordination with multiple agents
- Generating novel patterns or categories
- Creativity

⁵ *Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages*, McKinsey Global Institute, November 28, 2017.

- Logical reasoning and problem solving
- Output articulation or display
- Sensory perception

Solution development

We also updated our estimates for how long it would take to develop a solution that could integrate automation technologies once technical feasibility was established, based on a series of expert interviews. The original estimates were developed by examining the development time and technical capabilities for more than 100 previously developed automation solutions, including hardware and software solutions.

Economic feasibility

Once a solution is developed, we assume activities will start to be automated when the cost of the solution falls below the level of wages paid to a human to do that activity. For this calculation, we accounted for the evolution of solution costs as well as the evolution of wages.

Solution cost evolution

Based on the capability requirements, solutions are classified into two categories: hardware and software. If a solution requires sensory perception, fine motor skills and dexterity, gross motor skills, or mobility, it is classified as a hardware solution. Otherwise, it is classified as a software solution. For a given solution, the initial cost is estimated as a percentage of the highest hourly wage for the corresponding activity across all the countries we modeled. We estimated the initial cost by examining several examples of solutions developed using different mixes of hardware and software. Based on our research, most software solutions have relatively low initial costs as a percentage of the human labor cost. Some solutions that require a combination of both software and hardware components have a higher initial cost. To be conservative, we excluded certain solutions with advantages that could be derived only from very specific scenarios or that include noneconomic benefits such as increased quality and efficiency or decreased error rates. The range of initial solution costs we modeled is 20 to 70 percent of the highest hourly wage for the corresponding activity for hardware, and 0 to 20 percent for software.

In our model, solution costs decrease as technology advances. Hardware solution costs decline by 16.0 percent per year, and software solution costs decline by 5.3 percent per year. We triangulate the consumer price index and supplier surveys to estimate the reduction in hardware solution costs. For this, we use computer software and accessories indexes to estimate software solution cost reduction. For consumer price inflation, we used consumer price index data for personal computers and peripheral equipment. For computer software and accessories, we use data from the US Bureau of Labor Statistics. For software solution costs, we use a survey of prices from the International Federation of Robotics. Further work could be done to refine the estimation; however, given software's low starting cost, annual reductions have little impact on final automation results.

Wage evolution

We model the wage evolution for each country in two stages. From 2023 to 2030, we apply country-level growth estimates for all countries and convert them into 2010 constant US dollars by dividing nominal GDP by the corresponding country-level price deflator (2010 base) and multiply the exchange rate to the dollar (2010 base). We then calculate a CAGR for each country from 2022 to 2030. For 2030 onward, we grouped the 47 countries for which we have data into two cohorts using a cutoff country-level annual wage, based on the wage distribution from 2022 to 2030. Countries within the same cohort grow at the same rate. We reclassify countries each year. As a country advances into the next cohort, the appropriate growth rate is applied.

Adoption and deployment

Adoption can start once automation solutions are economically feasible, but several factors can hinder or enable the timing and the pace of adoption. Solutions requiring different technologies have varying levels of ease of integration. It takes time to integrate capabilities into current technical platforms and combine them into an organic entity. Barriers also exist on the organization side. Human talent and organization structures might act as bottlenecks to implementation. Policies and regulations could also slow down or accelerate technology innovation and adoption. Finally, depending on their preferences, consumers might have varied levels of acceptance for automated solutions that could affect the pace of adoption.

To incorporate all these factors, we used the mathematics of the Bass diffusion model, a well-known and widely used function in forecasting, especially for new product sales forecasting and technology forecasting.

$$\frac{f(t)}{1-F(t)} = (p + qF(t))$$

$F(t)$ is the installed base fraction (that is, adoption of given technology or product) and $f(t)$ is the corresponding rate of change.

The function in our case also contains two key parameters: p parameter (the inherent tendency of consumers to adopt new technology), and q parameter (the tendency of consumers to adopt based on peer adoption). Parameters are estimated through ordinary least square regression. In the absence of data, p and q parameter values from meta analyses can be used if a saturation value is known or can be guessed.

We then simulated two scenarios for historic technology adoption curves. The technologies we used are stents, airbags, laparoscopic surgery, MRI, smartphones, TVs, antilock braking systems, online air booking, cellphones, color TVs, SX/EW leaching, personal computers, electronic stability control, instrument landing systems, dishwashers, and pacemakers. The fitted values of parameters p and q are consistent with other academic research. It takes about five years to reach 50 percent adoption in the earliest scenario and approximately 16 years in the latest scenario.

2. Work hours that could be automated

The impact of automation on work hours across different capability levels is estimated below for the technical automation potential and automation adoption:

Impact by occupation by scenario: Number of FTEs X % automation estimate by scenario

Impact by activity: Number of FTEs X time spent on activity per year⁶ X percent automation estimate by scenario

Impact by occupation group⁷: $\sum_{\text{occupation}}$ (Number of FTEs X percent automation estimate by scenario)

3. Impact of automation on productivity

We used GDP per FTE as the measure of productivity. We calculated automation output under different scenarios by multiplying the projected number of FTEs by the estimated automation adoption rate. To maintain consistency with other data sources, we made several additional assumptions. We considered only job activities that are available and well defined as of the date of this report. Also, to be conservative, we assumed that automation has a labor

⁶ FTEs whose job titles include these detailed work activities.

⁷ For occupations included within the occupation group.

substitution effect but no other performance gains. Finally, we assumed that labor replaced by automation will rejoin the workforce at the same level of productivity as today. We also assumed that additional output from automation will not decrease, even if the total number of FTEs declines as a result of demographic changes.

Under the assumptions outlined above, we calculated the GDP impact of automation adoption by country as follows:

Additional GDP impact of automation = FTE impact of automation X productivity

The additional GDP impact of automation is then added to 2022 GDP to estimate the productivity impact of automation.

Acknowledgments

The research underpinning this report was led by Michael Chui, an MGI partner in McKinsey's Bay Area office; Eric Hazan, a senior partner in the Paris office; Roger Roberts, a partner in the Bay Area office; Alex Singla, a senior partner in the Chicago office; Kate Smaje and Alex Sukharevsky, senior partners in the London office; Lareina Yee, a senior partner in the Bay Area office; and Rodney Zimmel, a senior partner in the New York office.

The project team included Dmitry Gafarov, Shivani Gupta, Dan Hababou, Leila Harouchi, Sonja Lindberg, Kerin Lo, Alexandre Pons, Alok Singh, Gurneet Singh Dandona, and Wilbur Wang.

This research benefited immensely from the expertise and perspectives of many McKinsey colleagues. Special thanks to Pedro Abreu, a principal data scientist in the London office; Begum Karaci Deniz, a consultant in the Bay Area office; Kweilin Ellingrud, a senior partner and McKinsey Global Institute director in Shanghai; John Larson, an expert associate partner in the Southern California office; Damian Lewandowski, an expert in the New York office; Guillaume Lurenbaum, an expert in the Paris office; Matej Macak, a partner in the London office; and Marco Piccitto, a senior partner and McKinsey Global Institute director in Milan.

We also thank the following McKinsey colleagues: Rohit Agarwal, Steven Aronowitz, Arun Arora, Charles Atkins, Elia Berteletti, Onno Boer, Albert Bollard, Xavier Bosquet, Benjamin Braverman, Charles Carcenac, Sebastien Chaigne, Peter Crispeels, Santiago Comella-Dorda, Eleonore Depardon, Thierry Ethevenin, Neel Gandhi, Eric Goldberg, Liz Grennan, Vinay Gupta, Bryan Hancock, Lisa Harkness, Jake Hart, Heiko Heimes, Jeff Jacobs, Tarun Khurana, Malgorzata Kmicinska, Jan-Christoph Köstring, Andreas Kremer, Kathryn Kuhn, Jessica Lamb, Maxim Lampe, Swan Leroi, Richard Li, Dana Maor, Julien Mauhourat, Carolyn Pierce, Olivier Plantefevre, Kathryn Rathje, Emily Reasor, Werner Rehm, Steve Reis, Kelsey Robinson, Martin Rosendahl, Christoph Sandler, Saurab Sanghvi, Boudhayan Sen, Joanna Si, François Soubien, Eli Stein, Michele Tam, Robert Tas, Maribel Tejada, Georg Winkler, Jane Wong, and Romain Zilahi.

We are grateful to the following external advisers, who challenged our thinking and added new insights: Martin Neil Baily, senior fellow emeritus in economic studies at the Brookings Institution; Ethan Mollick, associate professor of management at the Wharton School of the University of Pennsylvania and academic director of Wharton Interactive; Éric Moulines, professor at École Polytechnique, Institut Polytechnique de Paris, co-scientific director of Hi! Paris, and member of the l'Académie des Sciences; and Gaël Richard, researcher and professor at Télécom Paris, Institut Polytechnique de Paris, and co-scientific director of Hi! Paris.

Additionally, we thank our McKinsey colleague Max Gleischman, global director of communications, reputation risk.

This report was edited by MGI senior editor Stephanie Strom and David DeLallo, executive editor at McKinsey Global Publishing. We also thank our colleagues David Batcheck, Tim Beacom, Nienke Beuwer, Chuck Burke, Amanda Covington, Ashley Grant, Cathy Gui, Vasudha Gupta, Diane Henry, Marion Obadia, Moira Pierce, Rebeca Robboy, Rachel Robinson, Katherine Shenton, Cindy Van Horne, and Nathan Wilson.

This research is independent and fact-based. None of it was commissioned or funded by any business, government, or other institutions, and we share it publicly free of charge. The research was entirely funded by the partners of McKinsey. While we engage multiple distinguished external advisers to contribute to our work, the analysis in this publication is ours alone and any errors are our own.

These McKinsey practices collaborated in the research and production of this report:

QuantumBlack, AI by McKinsey, is McKinsey's AI arm and helps companies transform using the power of technology, technical expertise, and industry experts. With thousands of practitioners at QuantumBlack (data engineers, data scientists, product managers, designers, software engineers) and McKinsey (industry and domain experts), we are working to solve the world's most important AI challenges. For more information, please visit McKinsey.com/capabilities/quantumblack.

McKinsey Digital drives transformation and builds businesses by bringing together the capabilities needed to help organizations grow and thrive in the digital age. We help our clients harness the power of data and artificial intelligence, modernize core technology and capitalize on new technology, optimize and automate operations, fuel digital growth, create stunning digital experiences, and build digital talent and culture. For further information, please visit McKinsey.com/capabilities/mckinsey-digital.

The McKinsey Technology Council brings together a global group of more than 100 scientists, entrepreneurs, researchers, and business leaders. We research, debate, inform, and advise, helping executives from all sectors navigate the fast-changing technology landscape.

The McKinsey Global Institute (MGI) was established in 1990. Our mission is to provide a fact base to aid decision making on the economic and business issues most critical to the world's companies and policy leaders. We benefit from the full range of McKinsey's regional, sectoral, and functional knowledge, skills, and expertise, but editorial direction and decisions are solely the responsibility of MGI directors and partners. We aim for independent and fact-based research and analysis. None of our work is commissioned or funded by any business, government, or other institution; we share our results publicly free of charge; and we are entirely funded by the partners of McKinsey. For further information about MGI and to download all reports for free, please visit McKinsey.com/mgi.

McKinsey Growth, Marketing & Sales delivers innovation at scale and growth transformation at speed to help solve the toughest challenges that stand in our clients' way. We bring the best of McKinsey—growth strategy, marketing and sales expertise, business building, and tech-enabled solutions—to help companies elevate growth. Our integrated approach is fueled by advanced analytics, such as growth mapping, predictive customer experience, pricing, performance marketing, and commercial growth. We put this actionable intelligence in our clients' hands to help them sustain growth. The result is a future-ready organization poised to accelerate growth. For further information, please visit McKinsey.com/capabilities/growth-marketing-and-sales.





McKinsey & Company
June 2023
Copyright © McKinsey & Company
www.McKinsey.com

 @McKinsey
 @McKinsey
 @McKinsey



NBER WORKING PAPER SERIES

GENERATIVE AI AT WORK

Erik Brynjolfsson
Danielle Li
Lindsey R. Raymond

Working Paper 31161
<http://www.nber.org/papers/w31161>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023, revised November 2023

We are grateful to Daron Acemoglu, David Autor, Amittai Axelrod, Eleanor Dillon, Zayd Enam, Luis Garicano, Alex Frankel, Sam Manning, Sendhil Mullainathan, Emma Pierson, Scott Stern, Ashesh Rambachan, John Van Reenen, Raffaella Sadun, Kathryn Shaw, Christopher Stanton, Sebastian Thrun, and various seminar participants for helpful comments and suggestions. We thank Max Feng for providing excellent research assistance and the Stanford Digital Economy Lab for funding. The content is solely the responsibility of the authors and does not necessarily represent the official views of Stanford University, MIT, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w31161>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Generative AI at Work
Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond
NBER Working Paper No. 31161
April 2023, revised November 2023
JEL No. D8,J24,M15,M51,O33

ABSTRACT

New AI tools have the potential to change the way workers perform and learn, but little is known about their impacts on the job. In this paper, we study the staggered introduction of a generative AI-based conversational assistant using data from 5,179 customer support agents. Access to the tool increases productivity, as measured by issues resolved per hour, by 14% on average, including a 34% improvement for novice and low-skilled workers but with minimal impact on experienced and highly skilled workers. We provide suggestive evidence that the AI model disseminates the best practices of more able workers and helps newer workers move down the experience curve. In addition, we find that AI assistance improves customer sentiment, increases employee retention, and may lead to worker learning. Our results suggest that access to generative AI can increase productivity, with large heterogeneity in effects across workers.

Erik Brynjolfsson
Stanford Digital Economy Laboratory
353 Jane Stanford Way, Office 136
Stanford, CA 94305
and NBER
erik.brynjolfsson@gmail.com

Lindsey R. Raymond
MIT Sloan School of Management
100 Main Street
E62-489
Cambridge, MA 02142
lindsey.r.raymond@gmail.com

Danielle Li
MIT Sloan School of Management
100 Main St, E62-484
Cambridge, MA 02142
and NBER
d_li@mit.edu

The emergence of generative artificial intelligence (AI) has attracted significant attention, but few studies have examined its economic impact. While various generative AI tools have performed well in laboratory settings, excitement about their potential has been tempered by concerns that these tools may be less effective in real-world settings, where they may encounter unfamiliar problems, face organizational resistance, or provide misleading information in a consequential environment (Peng et al., 2023a; Roose, 2023).

In this paper, we study the adoption of a generative AI tool that provides conversational guidance for customer support agents.¹ This is, to our knowledge, the first study of the impact of generative AI when deployed at scale in the workplace. We find that access to AI assistance increases the productivity of agents by 14%, as measured by the number of customer issues they are able to resolve per hour. In contrast to studies of prior waves of computerization, we find that these gains accrue disproportionately to less-experienced and lower-skill workers.² We argue that this occurs because generative AI systems work by capturing and disseminating the patterns of behavior that characterize the most productive agents, including knowledge that has eluded automation from earlier waves of computerization.

Computers and software have transformed the economy with their ability to perform certain tasks with far more precision, speed, and consistency than humans. To be effective, these systems typically require explicit and detailed instructions for how to transform inputs into outputs: when engineers write code to perform a task, they are codifying that task.³ Yet because many workplace activities—such as writing emails, analyzing data, or creating presentations—rely on tacit knowledge, they have so far defied automation (Polanyi, 1966; Autor, 2014).⁴

Machine learning (ML) algorithms work differently from traditional computer programs: instead of requiring explicit instructions to function, these systems infer instructions from examples. Given a training set of images, for instance, ML systems can learn to recognize specific individuals even though one cannot fully explain what physical features characterize a given person’s identity. This ability highlights a key, distinguishing aspect of ML systems: they can learn to perform tasks even

¹A note on terminology. There are many definitions of artificial intelligence and of intelligence itself—Legg et al. (2007) list over 70 of them. In this paper, we define “artificial intelligence” (AI) as an umbrella term that refers to systems that exhibit intelligent behavior, such as learning reasoning and problem-solving. “Machine learning” (ML) is a branch of AI that uses algorithms to learn from data, identify patterns, and make predictions or decisions without being explicitly programmed (Google, n.d.). Large language models (LLMs) and tools built around LLMs such as ChatGPT are an increasingly important application of machine learning. LLMs generate new content, making them a form of “generative AI.”

²We provide a discussion of this literature at the end of this section.

³By codify, we mean compiling a process into a formal, ordered routine.

⁴Tacit knowledge refers to the knowledge and skills that individuals possess but are unable to express explicitly. It is often intuitive and nonverbal, acquired through personal experiences, observations, and practice over time. Tacit knowledge is deeply ingrained in an individual’s behavior and can be difficult to transfer or convey to others through traditional methods such as training or manuals (Polanyi, 1966).

when no instructions exist—including tasks requiring tacit knowledge that could previously only be gained through lived experience (Polanyi, 1966; Autor, 2014; Brynjolfsson and Mitchell, 2017).⁵

In addition, ML systems are often trained on data from human workers, who naturally differ in their abilities. By seeing many examples of tasks—making sales pitches, driving a truck, or diagnosing a patient, to name a few—performed well and poorly, these models can implicitly learn what specific behaviors and characteristics set high-performing workers apart from their less effective counterparts. That is, not only are generative AI models capable of performing complex tasks, they might also be capable of capturing the skills that distinguish top workers. The use of ML tools may therefore expose lower-skill workers to new skills and lead to differential changes in productivity.

We study the impact of generative AI on productivity and worker experience in the customer service sector, an industry with one of the highest rates of AI adoption (Chui et al., 2021). We examine the staggered deployment of a chat assistant using data from 5,000 agents working for a Fortune 500 software firm that provides business process software. The tool we study is built on a recent version of the Generative Pre-trained Transformer (GPT) family of large language models developed by OpenAI (OpenAI, 2023). It monitors customer chats and provides agents with real-time suggestions for how to respond. It is designed to augment agents, who remain responsible for the conversation and are free to ignore its suggestions.

We have three sets of findings.

First, AI assistance increases worker productivity, resulting in a 14% increase in the number of chats that an agent successfully resolves per hour. This increase reflects shifts in three components of productivity: a decline in the time it takes an agent to handle an individual chat, an increase in the number of chats that an agent handles per hour (agents may handle multiple chats at once), and a small increase in the share of chats that are successfully resolved. The productivity impacts of AI assistance are highly uneven. We find that less-skilled and less-experienced workers improve significantly across all productivity measures we consider, including a 34% increase in the number of issues they are able to resolve per hour. Access to the AI tool helps newer agents move more quickly down the experience curve: treated agents with two months of tenure perform just as well as untreated agents with more than six months of tenure. In contrast, we find minimal impacts on the productivity of more-experienced or more-skilled workers. Indeed, we find evidence that AI assistance may decrease the quality of conversations by the most skilled agents. These results contrast, in spirit, with studies that find evidence of skill-biased technical change for earlier waves

⁵As Meijer (2018) puts it “where the Software 1.0 Engineer formally specifies their problem, carefully designs algorithms, composes systems out of subsystems or decomposes complex systems into smaller components, the Software 2.0 Engineer amasses training data and simply feeds it into an ML algorithm...”

of computer technology (Autor et al., 2003; Acemoglu and Restrepo, 2018; Bresnahan et al., 2002; Bartel et al., 2007).

Our second set of results investigates the mechanism underlying our main findings. We show that AI recommendations appear useful to workers: agents who follow recommendations more closely see larger gains in productivity, and adherence rates increase over time for all workers, particularly those who were initially more skeptical. We also find that engagement with AI recommendations can generate durable learning. Using data on software outages—periods in which the AI software fails to provide any suggestions—we show that workers see productivity gains relative to their pre-AI baseline even when recommendations are unavailable. These outage-period gains are more pronounced for workers who had more prior exposure to AI assistance or who had followed AI recommendations more closely. Finally, we analyze the text of agents’ chats and provide suggestive evidence that access to AI drives convergence in communication patterns: low-skill agents begin communicating more like high-skill agents.

Our third set of results focus on agents’ experience of work. Work in contact centers⁶ is often difficult. Agents are regularly exposed to hostile treatment from upset (and anonymous) customers, and because much work is outsourced, many agents work overnight shifts in order to service US business hours. AI assistance may help agents communicate more effectively but could also increase the likelihood that agents are perceived as mechanical or inauthentic. We show that access to AI assistance markedly improves how customers treat agents, as measured by the sentiment of their chat messages. We also find that customers are less likely to question the competence of agents by requesting to speak to a supervisor. These changes come alongside a substantial decrease in worker attrition, which is driven by the retention of newer workers.

Our overall findings show that access to generative AI can increase the productivity and retention of individual workers. We emphasize, however, that our paper is not designed to shed light on the aggregate employment or wage effects of generative AI tools. Firms may respond to increasing productivity among novice workers by hiring more of them, de-skilling positions, or seeking to develop more powerful AI systems that can replace lower-skill workers entirely. Unfortunately, our data do not allow us to observe changes in wages, overall labor demand, or the skill composition of workers hired for the job.

Our results also highlight the longer-term incentive challenges that AI systems bring. Top workers are generally not paid for their contributions to the training data that AI systems use to capture and disseminate their skills. Yet, without these contributions, AI systems may be less

⁶The term “contact center” updates the term “call center,” to reflect the fact that a growing proportion of customer service contacts no longer involve phone calls.

effective in learning to resolve new problems. Our work therefore raises questions about how workers should be compensated for the data they provide to AI systems.

Our paper is related to a large literature on the impact of technological adoption on worker productivity and the organization of work (e.g. Rosen, 1981; Autor et al., 1998; Athey and Stern, 2002; Bresnahan et al., 2002; Bartel et al., 2007; Acemoglu et al., 2007; Hoffman et al., 2017; Bloom et al., 2014; Michaels et al., 2014; Garicano and Rossi-Hansberg, 2015; Acemoglu and Restrepo, 2020; Felten et al., 2023). Many of these studies, particularly those focused on information technologies, find evidence that IT complements higher-skill workers (Akerman et al., 2015; Taniguchi and Yamada, 2022). Bartel et al. (2007) shows that firms that adopt IT tend to use more skilled labor and increase skill requirements for their workers. Acemoglu and Restrepo (2020) study the diffusion of robots and find that the negative employment effects of robots are most pronounced for workers in blue collar occupations and those with less than a college education.

There have been substantially fewer studies involving AI-based technologies, generative or not. Acemoglu et al. (2022); Zolas et al. (2020); Calvino and Fontanelli (2023) examine economy-wide data from the US and OECD and show that the adoption of AI tools is concentrated among large, young firms with relatively high productivity. So far, evidence on the productivity impacts of these technologies is mixed: for example, Acemoglu et al. (2022) finds no detectable relationship between investments in AI-specific tools, while Babina et al. (2022) finds evidence of a positive relationship between firms’ AI investments and their subsequent growth and valuations.⁷ These studies all caution that the productivity effects of AI technologies may be challenging to identify at the macro-level because AI-adopting firms differ substantially from non-adopters.

In this paper, we provide micro-level evidence on the adoption of a generative AI tool across thousands of workers working at a given firm and its subcontractors. Our work is more closely related to several other studies examining the impacts of generative AI in lab-like settings. Peng et al. (2023b) recruit software engineers for a specific coding task (writing an HTTP server in JavaScript) and show that those given access to GitHub Copilot complete this task twice as quickly. Similarly, Noy and Zhang (2023) conduct an online experiment showing that subjects given access to ChatGPT complete professional writing tasks more quickly. Choi and Schwarcz (2023) give law students access to AI assistance on a law school exam. Consistent with our findings, Noy and Zhang (2023), Choi and Schwarcz (2023) and Peng et al. (2023a) find that ChatGPT compresses the productivity distribution, with lower-skill workers benefiting the most. Our paper, however, is

⁷OECD (2023) reports that when surveyed firms are asked directly, 57% of employers in finance and 63% in manufacturing reported that AI positively impacted productivity and 80% of surveyed workers who work with AI report higher job performance.

the first to examine longer-term effects in a real-world workplace where we can also track patterns of learning, customer-side effects and changes in the experience of work.

1 Generative AI and Large Language Models

In recent years, the rapid pace of AI development and public release tools such as ChatGPT, GitHub Copilot, and DALL-E have attracted widespread attention, optimism, and alarm ([The White House, 2022](#)). These technologies are all examples of “generative AI,” a class of machine learning technologies that can generate new content—such as text, images, music, or video—by analyzing patterns in existing data. In this section, we provide background on generative AI as a technology and discuss its potential economic implications.

1.1 Technical Primer

This paper focuses on an important class of generative AI, large language models (LLMs). LLMs are neural network models designed to process sequential data ([Bubeck et al., 2023](#)). An LLM is trained by learning to predict the next word in a sequence, given what has come before, using a large corpus of text (such as Wikipedia, digitized books, or portions of the Internet). This knowledge of the statistical co-occurrence of words allows it to generate new text that is grammatically correct and semantically meaningful. Though “large language model” implies human language, the same techniques can be used to produce other forms of sequential data (“text”) such as protein sequences, audio, computer code, or chess moves ([Eloundou et al., 2023](#)).

Recent progress in generative AI has been driven by four factors: computing scale, earlier innovations in model architecture, the ability to “pre-train” using large amounts of unlabeled data and refinements in training techniques.⁸

First, the quality of LLMs is strongly dependent on scale: the amount of computing power used for training, the number of model parameters, and dataset size ([Kaplan et al., 2020](#)). Firms are increasingly devoting more resources to increasing this scale. The GPT-3 model included 175 billion parameters, was trained on 300 billion tokens, and generated approximately \$5 million dollars in computing costs alone; the GPT-4 model, meanwhile, is estimated to include 1.8 trillion parameters, trained on 13 trillion tokens, at a rumored computing-only cost of \$ 65 million ([Li, 2020](#); [Brown et al., 2020](#); [Patel and Wong, 2023](#))

⁸For a more detailed technical review of progress, see [Radford and Narasimhan \(2018\)](#); [Radford et al. \(2019\)](#); [Liu et al. \(2023\)](#); [Ouyang et al. \(2022\)](#).

In terms of model architecture, modern LLMs use two earlier key innovations: positional encoding and self-attention. Positional encodings keep track of the order in which a word occurs in a given input.⁹ Meanwhile, self-attention assigns importance weights to each word in the context of the entire input text. Together, this approach enables models to capture long-range semantic relationships within an input text, even when that text is broken up into smaller segments and processed in parallel (Vaswani et al., 2017; Bahdanau et al., 2015).

Next, LLMs can be pre-trained on large amounts of unlabeled data from sources such as Reddit or Wikipedia. Because unlabeled data is much more prevalent than labeled data, LLMs can learn about natural language on a much larger training corpus (Brown et al., 2020). By seeing, for instance, that the word “yellow” is more likely to be observed with “banana” or “sun” or “rubber duckie,” the model can learn about semantic and grammatical relationships even without explicit guidance (Radford and Narasimhan, 2018). The resulting model can be used in multiple applications because its training is not specific to a particular set of tasks.

Finally, general-purpose LLMs can be further “fine-tuned” to generate output that matches the priorities of any specific setting (Ouyang et al., 2022; Liu et al., 2023). For example, a model trained to generate social media content would benefit from receiving labeled data that contain not just the content of a post or tweet, but also information on the amount of user engagement it received. Similarly, an LLM may generate several potential responses to a given query, but some of them may be factually incorrect or contain toxic language. To discipline this model, human evaluators can rank these outputs to train a reward function that prioritizes desirable responses. These types of refinements can significantly improve model quality by making a general-purpose model better suited to its specific application (Ouyang et al., 2022).

Together, these innovations have generated meaningful improvements in model performance. The Generative Pre-trained Transformer (GPT) family of models, in particular, has attracted considerable media attention for their rapidly expanding capabilities.¹⁰

1.2 The Economic Impacts of Generative AI

Computers have historically excelled at executing pre-programmed instructions, making them particularly effective at tasks that can be reduced to explicit rules (Autor, 2014). Consequently, computerization has disproportionately reduced demand for workers performing “routine” tasks such as data entry, bookkeeping, and assembly line work, reducing wages in these jobs (Acemoglu and Autor, 2011). At the same time, computerization has also increased the demand for workers who

⁹For instance, a model would keep track of “the, 1” instead of only “the” (if “the” was the first word in the sentence).

¹⁰For instance, GPT-4 has recently been shown to outperform humans in taking the US legal bar exam (Liu et al., 2023; Bubeck et al., 2023; OpenAI, 2023).

possess complementary skills such as programming, data analysis, and research. Together, these changes have contributed to increasing wage inequality in the United States and have been linked to a variety of organizational changes (Katz and Murphy, 1992; Autor et al., 2003; Michaels et al., 2014; Bresnahan et al., 2002; Baker and Hubbard, 2003; OECD, 2023).

In contrast, generative AI tools do not require explicit instructions to perform tasks. If asked to write an email denying an employee a raise, generative AI tools will likely respond with a professional and conciliatory note. This occurs because the model will have seen many examples of workplace communication in which requests are declined in this manner. Importantly, the model produces such an output even though no programmer has explicitly specified what tone would be appropriate for what context, nor even defined what a tone like “professional” or “conciliatory” means. Indeed, the ability to behave “appropriately” is one that cannot be fully articulated even by those who possess it. Rather, people learn to do so from experience and apply unconscious rules in the process. This type of “tacit knowledge” underlies most tasks humans perform, both in and out of the workplace (Polanyi, 1966; Autor, 2014).

The fact that generative AI models display such skills suggests that they can acquire tacit knowledge that is embedded in the training examples they encounter. This ability expands the types of tasks that computers may be capable of performing to include non-routine tasks that rely on judgment and experience. For example, Github Copilot, an AI tool that generates code suggestions for programmers, has achieved impressive performance on technical coding questions and, if asked, can provide natural language explanations of how the code it produces works (Nguyen and Nadi, 2022; Zhao, 2023). Meanwhile, “AI-assistant” services such as Claude can be used to produce convincing business case analyses, including reading and interpreting financial statements and offering strategic assessments. Because many of these tasks—coding, financial analysis, etc.—are currently performed by workers who have either been insulated or benefited from prior waves of technology adoption, the expansion of generative AI has the potential to shift the relationship between technology, labor productivity, and inequality (The White House, 2022).

Generative AI tools can not only expand the types of tasks that machines can perform, they may also reveal valuable information about how the most productive human workers differ from others. This is because the ML models underlying generative AI systems are commonly trained on data generated by human workers and, consequently, encounter many examples of people performing tasks both well and poorly. In learning to predict good outcomes on such data, ML models may implicitly identify characteristics or patterns of behavior that distinguish high and low performers, including subtleties rooted in tacit knowledge. Generative AI systems then take this knowledge and use it to produce new behaviors that embody what top performers might do. This ability could be

used in different ways: firms may choose to replace lower-skill workers with AI-based tools, such tools could be used to demonstrate best practices to help lower-skill workers improve or help less experienced workers get up to speed more quickly. In either case, generative AI tools may have differential impacts by worker ability, even amongst workers performing the same tasks.

Despite their potential, generative AI tools face significant challenges in real-world applications. At a technical level, popular LLM-based tools, such as ChatGPT, have been shown to produce false or misleading information in unpredictable ways, generating concern about their ability to be reliable in high-stakes situations. Second, while LLM models often perform well on specific tasks in the lab (OpenAI, 2023; Peng et al., 2023b; Noy and Zhang, 2023), the types of problem that workers encounter in real-world settings are likely to be broader and less predictable. This raises concerns both about whether AI tools will be able to provide accurate assistance in every circumstance and—perhaps more importantly—about whether workers will be able to distinguish cases where AI tools are effective from those where they are not. Finally, the efficacy of new technologies is likely to depend on how they interact with existing workplace structures. Promising technologies may have more limited effects in practice due to the need for complementary organizational investments, skill development, or business process redesign. Because generative AI technologies are only beginning to be used in the workplace, little is currently known about their impacts.

2 Our Setting: LLMs for Customer Support

2.1 Customer Support and Generative AI

We study the impact of generative AI in the customer service industry, an area with one of the highest surveyed rates of AI adoption.¹¹ Customer support interactions are important for maintaining a company’s reputation and building strong customer relationships, yet, as in many industries, there is substantial variation in worker productivity (Berg et al., 2018; Syverson, 2011).

Newer workers are also often less productive and require significant training. At the same time, turnover is high: industry estimates suggest that 60% of agents in contact centers leave each year, costing firms \$10,000 to \$20,000 dollars per agent (Buesing et al., 2020; Gretz and Jacobson, 2018). To address these workforce challenges, the average supervisor spends at least 20 hours per week coaching agents with lower performance (Berg et al., 2018). Faced with variable productivity, high turnover, and high training costs, firms are increasingly turning to AI tools (Chui et al., 2021).

¹¹For instance, of the businesses that report using AI, 22% use AI in their customer service centers (Chui et al., 2021).

At a technical level, customer support is well-suited for current generative AI tools. From an AI’s perspective, customer-agent conversations can be thought of as a series of pattern-matching problems in which one is looking for an optimal sequence of actions. When confronted with an issue such as “I can’t login,” an AI/agent must identify which types of underlying problems are most likely to lead a customer to be unable to log in and think about which solutions typically resolve these problems (“Can you check that caps lock is not on?”). At the same time, they must be attuned to a customer’s emotional response, making sure to use language that increases the likelihood that a customer will respond positively (“that wasn’t stupid of you at all! I always forget to check that too!”). Because customer service conversations are widely recorded and digitized, pre-trained LLMs can be fine-tuned for customer service using many examples of both successfully and unsuccessfully resolved conversations.

Customer service is also a setting where there is high variability in the abilities of individual agents. For example, top-performing agents are often more effective at diagnosing the underlying technical issue given a customer’s problem description. These workers often ask more questions before settling on a diagnosis of the problem; this takes longer initially, but reduces the likelihood that agents waste time trying to resolve the wrong problem. Such differences in agent behavior can often be inferred from the large amounts of training data that customer-service-specific AI models have access to. As a result, customer service is also a setting in which generative AI models can potentially encode some of the “best practices” that top-performing agents use.

In the remainder of this section, we provide details about the firm we study and the AI tool they adopt.

2.2 Data Firm Background

We work with a company that provides AI-based customer service support software (hereafter, the “AI firm”) to study the deployment of their tool at one of their client firms, (hereafter, the “data firm”).

Our data firm is a Fortune 500 enterprise software company that specializes in business process software for small and medium-sized businesses in the United States. It employs a variety of chat-based technical support agents, both directly and through third-party firms. The majority of agents in our sample work from offices located in the Philippines, with a smaller group working in the United States and in other countries. Across locations, agents are engaged in a fairly uniform job: answering technical support questions from US-based small business owners.

Chats are randomly assigned, and support sessions are relatively lengthy, averaging 40 minutes, with much of the conversation spent trying to diagnose the underlying technical problem. The job

requires a combination of detailed product knowledge, problem solving skills, and the ability to deal with frustrated customers.

Our firm measures productivity using three metrics that are standard in the customer service industry: “average handle time,” the average time an agent takes to finish a chat; “resolution rate,” the share of conversations that the agent successfully resolves; and “net promoter score,” (customer satisfaction), which is calculated by randomly surveying customers after a chat and calculating the percentage of customers who would recommend an agent minus the percentage who would not. A productive agent is able to field customer chats quickly while maintaining a high resolution rate and net promoter score.

Across locations, agents are organized into teams with a manager who provides feedback and training to agents. Once a week, managers hold one-on-one feedback sessions with each agent. For example, a manager might share the solution to a new software bug, explain the implication of a tax change, or suggest how to better manage customer frustration with technical issues. Agents work individually, and the quality of their output does not directly affect others. Agents are paid an hourly wage and bonuses based on their performance relative to other agents.

2.3 AI System Design

The AI system we study combines a recent version of GPT with additional ML algorithms specifically fine-tuned to focus on customer service interactions. The system is further trained on a large set of customer-agent conversations that have been labeled with a variety of outcomes and characteristics: whether the call was successfully resolved, how long it took to handle the call, and whether the agent in charge of the call is considered a “top” performer by the data firm. The AI firm then uses these data to look for conversational patterns that are most predictive of call resolution and handle time.

The AI firm further trains its model using a process similar in spirit to [Ouyang et al. \(2022\)](#) to prioritize agent responses that express empathy, provide appropriate technical documentation, and limit unprofessional language. This additional training mitigates some of the concerns associated with relying on LLMs to generate text.

Once deployed, the AI system generates two main types of output: 1) real-time suggestions for how agents should respond to customers and 2) links to the data firm’s internal documentation for relevant technical issues. In both cases, recommendations are based on a history of the conversation.¹²

¹²For example, the correct response when a customer says “I can’t track my employee’s hours during business trips” depends on what version of the data firm’s software the customer uses. Suppose that the customer has previously mentioned that they are using the premium version. In that case, they should have access to remote mobile device

Figure 1 illustrates an example of AI assistance. In the chat window (Panel A), Alex, the customer, describes their problem to the agent. Here, the AI assistant generates two suggested responses (Panel B). In this example, it has learned that phrases like “I can definitely assist you with this!” and “Happy to help you get this fixed asap” are associated with positive outcomes. Panel A of Appendix Figure A.1 shows an example of a technical recommendation from the AI system, which occurs when it recommends a link to the data firm’s internal technical documentation.

Importantly, the AI system we study is designed to augment, rather than replace, human agents. The output is shown only to the agent, who has full discretion over whether to incorporate (fully or partially) the AI suggestions. This reduces the likelihood that off-topic or incorrect outputs make their way into customer conversations. Furthermore, the system does not provide suggestions when it has insufficient training data for that situation. In these situations, the agent must respond on their own.

3 Deployment, Data, and Empirical Strategy

3.1 AI Model Deployment

The AI assistant we study was gradually rolled out at the agent level after an initial seven-week randomized pilot featuring 50 agents.¹³ The deployment was largely uniform across both the data firm’s own customer service agents and its outsourced agents. Appendix Figure A.2 documents the progression of deployment among agents who are eventually treated. The bulk of the adoption occurs between November 2020 and February 2021.

3.2 Summary Statistics

Table 1 provides details on sample characteristics, divided into three groups: agents who are never given access to the AI tool during our sample period (“never treated”), pre-AI observations for those who are eventually given access (“treated, pre”), and post-AI observations (“treated, post”). In total, we observe the conversation text and outcomes associated with 3 million chats by 5,179 agents. Within this, we observe 1.2 million chats by 1,636 agents in the post-AI period. Most agents in our sample, 89%, are located outside of the United States, primarily in the Philippines. For each agent, we observe their assigned manager, tenure, geographic location, and firm information.

timekeeping, meaning that the support agents need to diagnose and resolve a technical issue that prevents the software from working. If, however, the customer stated that they are using the standard version, then the correct solution is for the customer to upgrade to the premium version in order to access this feature. For more on context tracking, see, for instance, [Dunn et al. \(2021\)](#).

¹³Data from the RCT is included as part of our primary analysis but is not analyzed separately because of its small sample size.

To examine the impacts of this deployment, we construct several key variables, all aggregated to the agent-month level, which is our primary level of analysis.

Our primary measure of productivity is resolutions per hour (RPH), the number of chats a worker is able to successfully resolve per hour. We consider this measure to be the most effective summary of a worker’s productivity at the firm. An agent’s RPH is determined by several factors: the average time it takes an agent to complete a conversation, the number of conversations they are able to handle per hour (accounting for multiple simultaneous conversations), and the share of conversations that are successfully resolved. We measure these individually as, respectively, average handle time (AHT), chats per hour (CPH), and resolution rate (RR). In addition, we also observe a measure of customer satisfaction through an agent’s net promoter score (NPS), which is collected by the firm from post-call customer surveys.

We observe these measures for different numbers of agents. In particular, we are able to reconstruct measures of average handle time and chats per hour from our chat level data. We therefore observe AHT and CPH measures for all agents in our sample. Measures that involve an understanding of call quality—resolution rates, and customer satisfaction—are provided at the agent-month level by our data firm. Because our data firm outsources most of its customer service functions, it does not have direct control over this information, which is kept by subcontracted firms. As a result, we observe resolution rates and net promoter scores for a subset of agents in our data. This, in turn, means that we only observe our omnibus productivity measure—resolutions per hour—for this smaller subset.

Figure 2 plots the raw distributions of our outcomes for each of the never, pre-, and post-treatment subgroups. Several of our main results are readily visible in these raw data. In Panels A through D, we see that post-treatment agents do better along a range of outcomes, relative to both never-treated agents and pre-treatment agents. In Panel E, we see no discernible differences in surveyed customer satisfaction among treated and non-treated groups.

Focusing on our main productivity measure, Panel A of Figure 2 and Table 1 show that never-treated agents resolve an average of 1.7 chats per hour, whereas post-treatment agents resolve 2.5 chats per hour. Some of this difference may be due to differences in the initial section: treated agents have higher resolutions per hour prior to AI model deployment (2.0 chats) relative to never-treated agents (1.7). This same pattern appears for chats per hour (Panel C) and resolution rates (Panel D): while ever-treated agents appear to be stronger performers at the outset than agents who are never treated, post-treatment agents perform substantially better. When looking instead at average handle times (Panel B), we see a starker pattern: pre-treatment and never-treated agents have similar distributions of average handle times, centered at 40 minutes, but post-treatment agents

have a lower average handle time of 35 minutes. These figures, of course, reflect raw differences that do not account for potential confounding factors such as differences in agent experience or differences in selection into treatment. In the next section, we will more precisely attribute these raw differences to the impact of AI model deployment.

3.3 Empirical Strategy

We isolate the causal impact of access to AI recommendations using a standard difference-in-differences regression:

$$y_{it} = \delta_t + \alpha_i + \beta AI_{it} + \gamma X_{it} + \epsilon_{it} \quad (1)$$

Our outcome variables y_{it} capture various measures of productivity for agent i in year-month t , as outlined earlier. Because workers often work only for a portion of the year, we include only year-month observations for an agent who is actively employed (e.g. assigned to chats). Our main variable of interest is AI_{it} , an indicator equal to one if agent i has access to AI recommendations at time t . All regressions include year-month fixed effects δ_t to control for common, time-varying factors such as tax season or the end of the business quarter. In our preferred specification, we also include controls for time-invariant agent-level fixed effects α_i and time-varying agent tenure. Standard errors are clustered at the agent level.

A rapidly growing literature has shown that two-way fixed effects regressions deliver consistent estimates only with strong assumptions about the homogeneity of treatment effects, and may be biased when treatment effects vary over time or by adoption cohort (Cengiz et al., 2019; de Chaisemartin and D’Haultfœuille, 2020; Sun and Abraham, 2021; Goodman-Bacon, 2021; Callaway and Sant’Anna, 2021; Borusyak et al., 2022). For example, workers may take time to adjust to using the AI system, in which case its impact in the first month may be smaller. Alternatively, the onboarding of later cohorts of agents may be smoother, so that their treatment effects may be larger.

We study the dynamics of treatment effects using the interaction weighted (IW) estimator proposed in Sun and Abraham (2021). Sun and Abraham (2021) show that this estimator is consistent assuming parallel trends, no anticipatory behavior, and cohort-specific treatment effects that follow the same dynamic profile.¹⁴ In the appendix, we show that both our main differences-in-differences and event study estimates are similar using robust estimators introduced in de Chaisemartin and D’Haultfœuille (2020), Borusyak et al. (2022), Callaway and Sant’Anna (2021), and Sun and Abraham (2021), as well as using traditional two-way fixed effects OLS.

¹⁴This last assumption means that treatment effects are allowed to vary over event-time and that average treatment effects can vary across adoption-cohorts (because even if they follow the same event-time profile, we observe different cohorts for different periods of event-time).

4 Main Results

4.1 Productivity Metrics

Table 2 examines the impact of the deployment of the AI model on our primary measure of productivity, resolutions per hour, using a standard two-way fixed effects model. In Column 1, we show that, controlling for time and location fixed effects, access to AI recommendations increases resolutions per hour by 0.47 chats, up 22.2% from an average of 2.12. In Column 2, we include fixed effects for individual agents to account for potential differences between treated and untreated agents. In Column 3, we include additional controls for the time-varying agent tenure. As we add controls, our effects fall slightly, so that, with agent and tenure fixed effects, we find that the deployment of AI increases RPH by 0.30 chats or 13.8%. Columns 4 through 6 produce these same patterns and magnitudes for the log of RPH.

Appendix Table A.1 finds similar results using alternative difference-in-difference estimators introduced in Callaway and Sant’Anna (2021), Borusyak et al. (2022), de Chaisemartin and D’Haultfoeulle (2020), and Sun and Abraham (2021). Unlike traditional OLS, these estimators avoid comparing between newly treated and already treated units. In most cases, we find larger effects of AI assistance using these alternatives.

Figure 3 shows the accompanying IW event study estimates of Sun and Abraham (2021) for the impact of AI assistance on RPH, in levels and logs. For both outcomes, we find a substantial and immediate increase in productivity in the first month of deployment. This effect grows slightly in the second month and remains stable and persistent up to the end of our sample. Appendix Figure A.3 shows that this pattern can be seen using alternative event study estimators as well: Callaway and Sant’Anna (2021), Borusyak et al. (2022), de Chaisemartin and D’Haultfoeulle (2020), and traditional two-way fixed effects.

In Table 3, we report additional results using our preferred specification with year-month, agent, and agent tenure fixed effects. Column 1 documents a 3.8 minute decrease in the average duration of customer chats, a 9% decline from the baseline mean (shorter handle times are generally considered better). Next, Column 2 indicates a 0.37 unit increase in the number of chats that an agent can handle per hour. Relative to a baseline mean of 2.6, this represents an increase of roughly 14%. Unlike average handle time, chats per hour account for the possibility that agents may handle multiple chats simultaneously. The fact that we find a stronger effect on this outcome suggests that AI enables agents to both speed up chats and multitask more effectively.

Column 3 of Table 3 indicates a small 1.3 percentage point increase in chat resolution rates, significant at the 10% level. This effect is economically modest, given a high baseline resolution

rate of 82%; we interpret this as evidence that improvements in chat handling do not come at the expense of problem solving on average. Finally, Column 4 finds no economically significant change in customer satisfaction, as measured by net promoter scores: the coefficient is -0.13 percentage points and the mean is 79.6%. Columns 5 through 8 report these results for logged outcomes. Going forward, we will report our estimates in logs, for ease of interpretation.

Figure 4 presents the accompanying event studies for additional outcomes. We see immediate impacts on average handle time (Panel A) and chats per hour (Panel B), and relatively flat patterns for resolution rate (Panel C) and customer satisfaction (Panel D). We therefore interpret these findings as saying that, on average, AI assistance increases productivity without negatively impacting resolution rates and surveyed customer satisfaction.

4.2 Impacts by Agent Skill and Tenure

There is substantial debate about the distributional consequences of AI-based technologies on worker productivity. An extensive literature suggests earlier waves of information and communication technology (e.g., the Internet, computers, network-based communication) have complemented high-skill workers, increasing their productivity and labor demand and widening wage differentials. Generative AI tools, however, are based on machine learning tools that rely on looking for patterns associated with success. As discussed earlier, generative AI tools may have a different pattern of productivity consequences relative to earlier waves of technology adoption. In this section, we examine whether access to AI assistance has different impacts along two dimensions: worker skill and worker experience.

4.2.1 Pre-treatment Worker Skill

In Panel A of Figure 5, we consider how our estimated productivity effects differ by an agent’s pre-AI productivity. We divide agents into quintiles using a skill index based on their average call efficiency, resolution rate, and surveyed customer satisfaction in the quarter prior to the adoption of the AI system. These skill quintiles are defined within a firm-month. To isolate the impact of worker skill, we also control for worker tenure at AI deployment.

In Panel A, we show that the productivity impact of AI assistance is most pronounced for workers in the lowest skill quintile (leftmost side), who see a .29 log point or 34% increase in resolutions per hour. In contrast, AI assistance does not lead to any productivity increase for the most skilled workers (rightmost side).

In Figure 6 we show that less-skilled agents consistently see the largest gains across our other outcomes. For the highest-skilled workers, we find mixed results: a zero effect on average handle time

(Panel A), a positive effect for chats per hour (Panel B), and, interestingly, a small but statistically significant *decreases* in resolution rates and customer satisfaction (Panels C and D). These results are consistent with the idea that generative AI tools may function by exposing lower-skill workers to the best practices of higher-skill workers. Lower-skill workers benefit because AI assistance provides them with new solutions, whereas the best performers may see little benefit from being exposed to their own best practices. Indeed, the fact that we find negative effects along measures of chat quality—resolution rate and customer satisfaction—suggests that AI recommendations may distract top performers, or lead them to choose the faster option (following suggestions) rather than taking the time to come up with their own responses.

4.2.2 Pre-treatment Worker Experience

Next, we repeat our previous analysis for agent tenure. To do so, we divide agents into five groups based on their tenure at the time the AI model is introduced. Some agents have less than a month of tenure when they receive AI access, while others have more than a year of experience. To isolate the impact of worker tenure, we control for worker skill when given access to the AI.

In Panel B of Figure 5, we see a clear, monotonic pattern in which the least experienced agents see the greatest gains in resolutions per hour. Agents with less than 1 month of tenure improve their resolutions per hour by .38 log points or 46% improvement (relative to agents of the same tenure who do not have access to AI assistance). In contrast, we see no effect for agents with more than a year of tenure.

In Figure 7, we show the same patterns for other outcomes. In Panels A and B, we see that AI assistance generates large gains in call handling efficiency, measured by average handle times and chats per hour, respectively, among the newest workers. In Panels C and D, we find positive impacts of AI assistance on chat quality, as measured by resolution rates and customer satisfaction, respectively. For the most experienced workers, we see modest positive effects for average handle time (Panel A), positive but statistically insignificant effects on chats per hour (Panel B), and small but statistically significant negative effects for measures of call quality and customer satisfaction (Panels C and D).

4.2.3 Moving Down the Experience Curve

To further explore how AI assistance impacts newer workers, we examine how worker productivity evolves on the job.¹⁵ In Figure 8, we plot productivity variables by agent tenure for three distinct

¹⁵We avoid the term “learning curve” because we cannot distinguish if workers are learning or merely following recommendations.

groups: agents who never receive access to the AI model (“never treated”), those who have access from the time they join the firm (“always treated”), and those who receive access in their fifth month with the firm (“treated 5 mo.”).

We see that all agents begin with around 2.0 resolutions per hour. Workers who are never treated (blue line) slowly improve their productivity with experience, reaching approximately 2.5 resolutions per hour 8 to 10 months later. In contrast, workers who always have access to AI assistance (red line) increase their productivity to 2.5 resolutions per hour after only two months and continue to improve until they are resolving more than 3 chats per hour after five months of tenure.¹⁶ Comparing just these two groups suggests that access to AI recommendations helps workers move more quickly down the experience curve.

The final group in Panel A tracks workers who begin their tenure with the firm without access to AI assistance, but who receive access after five months on the job (green line). These workers improve slowly in the same way as never-treated workers for the first five months of their tenure. Starting in month five, however, these workers gain access and we see their productivity rapidly increase following the same trajectory as the always-treated agents. In Appendix Figure A.4, we plot these curves for other outcomes. We see clear evidence that the experience curve for always-treated agents is steeper for handle time, chats per hour, and resolution rates (Panels A through C). Panel D follows a similar but noisier pattern for customer satisfaction.

Together, these results indicate that access to AI helps new agents move more quickly down the experience curve. Across many of the outcomes in Figure 8, agents with two months of tenure and access to AI assistance perform as well as or better than agents with more than six months of tenure who do not have access.

5 Adherence, Learning, and Conversational Change

In this section, we conduct a variety of analyses aimed at better understanding the mechanisms behind our main results.

First, we examine how workers engage with AI recommendations. We show that workers are selective about the recommendations they adopt, following the recommendations 35% on average. We find that the returns to AI assistance are highest for workers who choose to follow recommendations. Consistent with a story in which workers find AI recommendations helpful, we show that adherence rates increase over time for all workers, especially among older workers: by the end of our sample, we see similar adherence rates across worker tenure and skill.

¹⁶Our sample ends here because we have very few observations more than five months after treatment.

Second, we explore whether AI-assistance helps workers learn. Using information on software outages in which AI assistance is temporarily unavailable, we provide evidence that exposure to AI leads to durable changes in worker skills. We find that workers exposed to AI recommendations continue to perform better during outages, and this effect is greater after more exposure and for agents who more closely follow AI recommendations when the software is working.

Lastly, using text-based analysis of chat records themselves, we provide suggestive evidence that AI assistance changes the content of agents’ communication. We document within-agent changes in communication following AI deployment, with larger changes for lower-skill workers. Across-person, we show that these changes increase the similarity of communication patterns between low- and high-skill agents. These results are consistent with AI recommendations leading lower-skill workers to communicate more like high-skill workers.

Taken together, our results suggest that examining and following AI recommendations helps workers—particularly lower-skilled workers—learn to adopt best practices gathered from higher-skill and more experienced agents.

5.1 Adherence to AI recommendations

The AI tool we study makes suggestions, but agents are ultimately responsible for what they say to the customer. In our main results, we estimate how access to the AI tool impacts outcomes regardless of how frequently agents follow its recommendations. Here, we examine how closely agents adhere to AI recommendations, and document the association between adherence and returns to adoption.

We measure “adherence” starting at the chat level, using the share of AI recommendations that each agent follows. Our AI firm codes agents as having adhered to a recommendation if they either click to copy the suggested AI text or if they self-input something very similar. We take this chat-level measure and aggregate it to the agent-month level.

Panel A of Figure 9 shows the distribution of average agent-month-level adherence for our post-AI sample, weighted by the log number of AI recommendations provided to that agent in that month. The average adherence rate is 38%, with an interquartile range of 23% to 50%: agents frequently ignore recommendations. In fact, the share of recommendations followed is similar to the share of other publicly reported numbers for generative AI tools; a study of GitHub Copilot reports that individual developers use 27% to 46% of code recommendations (Zhao, 2023). Such behavior may be appropriate, given that AI models may make incorrect or irrelevant suggestions.

Panel B of Figure 9 shows that *returns* to AI model deployment are higher when agents actually follow recommendations. To show this, we divide agents into quintiles based on the percent of AI recommendations they follow in the first month of AI access and separately estimate the impact of

AI assistance for each group. These estimates control for year-month and agent fixed effects as in Column 5 of Table 2.

We find a steady and monotonic increase in returns by agent adherence: among agents in the lowest quintile, we still see a 10% gain in productivity, but for agents in the highest quintile, the estimated impact is over twice as high, close to 25%. Appendix Figure A.5 shows the results for our other four outcome measures. The positive correlation between adherence and returns holds most strongly for average handle time (Panel A) and chats per hour (Panel B), and more noisily for resolution rate (Panel C) and customer satisfaction (Panel D).

Our results are consistent with there being a treatment effect of following AI recommendations on productivity. We note, however, that this relationship could also be driven by other factors: selection (agents who choose to adhere are more productive for other reasons); or selection on gains (agents who follow recommendations are those with the greatest returns). To further explore this, we consider worker’s revealed preference: do they continue to follow AI recommendations over time? If our results were driven purely by selection, we would expect workers with low adherence to continue having low adherence, since it was optimal for them to do so.

Figure 10 plots the evolution of AI adherence over time, for various categories of agents. Panel A begins by considering agents who differ in their initial AI compliance, which we categorize based on terciles of AI adherence in the first month of model deployment. Here, we see that compliance either stays stable or grows over time. The most initially compliant agents continue to comply at the same rates (just above 50%). Less initially compliant agents increase their compliance over time: those in the bottom tercile initially follow recommendations less than 20% of the time but, by month five, their compliance rates have increased by over 50%, to just over half of the time. Next, Panel B divides workers up by tenure at the time of AI deployment. More senior workers are initially less likely to follow AI recommendations: 30% for those with more than a year of tenure compared to 37% for those with less than three months of tenure. Over time, however, all workers increase their adherence, with more senior workers doing so faster so that the groups converge five months after deployment. In Panel C, we show the same analysis by worker skill at AI deployment. Here, we see that compliance rates are similar across skill groups, and all groups increase their compliance over time.

The results in Figure 10 are consistent with agents—particularly those who are initially more skeptical—coming to value AI recommendations over time. An alternative hypothesis, however, is that agents who dislike working with AI assistance exit the firm at higher rates. In Appendix Figure A.6 we repeat the analysis above, focusing on within-agent changes in adherence (that is, adherence rates residualized by agent fixed effects). Our within-agent results follow a similar pattern: all

workers increase adherence over time, and these increases appear largest for workers who were initially the least compliant and workers who were the most senior. This suggests that increases in adherence over time are not driven exclusively by selection.

5.2 Worker Learning

A key question raised by our findings so far is whether these improvements in productivity and changes in communication patterns reflect durable changes in the human capital of workers or simply their growing reliance on AI assistance.

To study this, we examine how workers perform during periods in which they are not able to access AI-recommendations due to technical issues at the AI firm. Outages occur occasionally in our data and can last anywhere from a few minutes to a few hours. During an outage, the system fails to provide recommendations to some, but not necessarily all, workers. For example, outages may affect agents who log into their computers after the system crashes, but not agents working at the same time who had signed in earlier. They may also affect workers using one physical server but not another. Our AI firm tracks the most significant outages in order to perform technical reviews of what went wrong. We compile these system reports to identify periods in which a significant fraction of chats are impacted by outages.

Appendix Figure A.7 shows an example of such an outage, which occurred on September 10, 2020. The y -axis plots the share of post-treatment chats (e.g. those occurring after the AI system has been deployed for a given agent) for which the AI software does not provide any suggestions, aggregated to the hour level. The x -axis tracks hours in days leading up to and following the outage event (hours with fewer than 15 post-treatment chats are plotted as zeros for figure clarity). During non-outage periods, the share of chats without AI recommendations is typically 30-40%. This reflects the fact that the AI system does not generate recommendations in response to all messages, even when it is functioning properly. Because many chats are short, it is common to see chats end without the AI system intervening. On the morning of September 10th, however, we see a notable spike in the number of chats without recommendations, increasing to almost 100%. Records from our AI firm indicate that this outage was caused by a software engineer running a load test that crashed the system.

Figure 11 examines the impact of access to the AI system for chats that occur during and outside these outage periods. Whereas our main event study regressions are at the worker-month level, these are at the chat level, in order to more precisely compare conversations that occurred during outage periods, versus those that did not. Panel A considers the impact of the introduction of AI assistance on chat duration (shorter is more efficient), using only post-adoption periods in

which no outages are reported. Consistent with our main results, we see an immediate decline in the duration of individual chats by approximately 10% to 15%.

In Panel B, we use the same pre-treatment observations, but now restrict to post-adoption periods that are impacted by large outages. We first note that our estimates are noisy and their magnitude appears larger than for non-outage periods (15% to 25% declines in chat duration). Because AI outages are rare and not necessarily random, this may reflect differences in the types of chats that are seen during outage periods than during non-outage periods. However, focusing on the size of estimated effects over time, an interesting pattern emerges. Rather than declining immediately post-adoption and staying largely stable as we see in Panel A for non-outage periods, Panel B shows that the benefit of exposure to AI assistance increases with time during outage periods. That is, if an outage occurs one month after AI adoption, workers do not handle the chat much more quickly than their pre-adoption baseline. Yet, if an outage occurs after three months of exposure to AI recommendations, workers handle the chat faster—even though they are not receiving direct AI assistance.

In Figure 12, we split our main outage event studies by worker’s initial AI adherence, as described in Section 5.1. Panel A shows that workers with high initial AI adherence see large and fast declines in chat processing times (relative to their pre-adoption baseline), even during outages. Panel B, in contrast, shows no such impact for workers who tend to deviate from AI recommendations: they see no improvement in chat times during outage periods, even after many months of AI access. These findings suggest that workers learn more by actively using AI suggestions.

Together, these results suggest that generative AI tools can help workers develop durable skills. Prior to the deployment of AI-assistance, agents only received training from managers during brief weekly coaching sessions. During these sessions, managers would go through several conversations from the past week and advise the worker on how they might have handled certain conversations better. However, by necessity, managers can only provide feedback on a small fraction of the conversations that an agent conducts. Moreover, because managers are often pressed for time and may lack training, they may simply point out weak metrics (“you need to lower your handle time”) rather than identifying strategies for how an agent could better approach a problem (“you need to ask more questions at the beginning to diagnose the issue better.”) This type of coaching is ineffective and can be counterproductive for employee engagement (Berg et al., 2018). In contrast, AI assistance provides workers with specific, actionable suggestions in real time. Our findings suggest that this can play a useful role in supplementing existing on-the-job training programs.

5.3 Conversational Change

Lastly, we consider how access to AI assistance influences how workers communicate. To capture an overall sense of the content of conversations, we begin by creating textual embeddings of agent-customer conversations. Textual embeddings take a given body of text and transform it into a high-dimensional vector that represents its “coordinates” in linguistic space. Two pieces of text will have more similar coordinates if they share a common meaning or style. The specific embedding given to a body of text will depend on the embedding model that is used. We form our text embeddings using all-MiniLM-L6-v2, an LLM that is specifically intended to capture and cluster semantic information to assess similarity across text (Hugging Face, 2023). Once we create an embedding for each conversation, we can compare the similarity of conversations by looking at the cosine similarities of their associated vectors; this common approach yields a score of 0 if two pieces of text are semantically orthogonal and a score of 1 if they have the same meaning (Koroteev, 2021). For context, the sentences “Can you help me with logging in?” and “Why is my login not working?” have a cosine similarity of 0.68 in our model.

Using this approach, we first show that AI assistance changes the content of what agents write to customer, rather than just typing the same things faster. Second, we explore how these patterns differ for high- and low-skill workers. We are particularly interested in understanding whether AI models can disseminate the behaviors of high performers. If this is the case, then we would expect AI assistance to lead lower-performing agents to write more like high-performers.

5.3.1 Within-worker changes in communication

We begin by examining how an agent’s communication evolves over time, before and after access to AI assistance. We begin by examining treated workers and comparing the similarity of their chats in each given event-time week to their chats from the month before AI deployment (week -4 to week -1). We exclude messages from the customer and focus only on agent-generated language. Panel A of Figure 13 plots the cosine similarity associated with these comparisons. We find that textual similarity to the pre-AI window is stable in the weeks leading up to the AI roll-out and drops immediately following AI deployment. That is, conversations 12 to 5 weeks before deployment are quite similar to conversations 4 to 1 week before, but conversations 0 to 12 weeks after are all less similar.

This drop in similarity is broadly inconsistent with the idea that AI assistance merely leads workers to type the same things but faster. If that were the case, we would expect call handle times to drop, but textual similarity to remain constant.

Next, Panel B of Figure 13 compares the magnitude of this pre- versus post-deployment change in textual content varies by pre-AI worker skill. We find that lower-skill agents (those in the bottom quintile of the pre-AI skill distribution) experience greater textual change after AI adoption, relative to top performers (those in the top quintile). Our results here control for firm-year-month fixed effects, which can account for seasonal changes in topics such as tax or payroll cycles, or new product rollouts. We also control for agent tenure fixed effects, which can account for the possibility that younger workers’ language may evolve more quickly independent of access to the AI model. Although we cannot control directly for the chat topic, we note that chat topics are randomly assigned to agents, so we would not expect differences in topics to vary systematically by agent skill. We interpret these results as providing suggestive evidence that AI deployment shifts the communication patterns of low-skill workers more than high-skill workers.

5.3.2 Across worker comparisons

Figure 14 considers whether individual level changes in communication lead low- and high-skill workers to sound more alike. To examine this, we plot the cosine similarity between high- and low-skill agents at specific moments in calendar time, separately for workers with (blue dots) and without (red diamonds) access to AI assistance. Among agents without AI access, we define high- and low-skill agents as those who are in the top or bottom quintile of our skill index for that month. Among agents with AI access, we define high- and low-skill agents based on whether they are in the top or bottom quintile of skill at the time of AI deployment.

Focusing on the blue dots, we see that the average textual similarity between high- and low-productivity workers is 0.55 among workers who do not have access to AI assistance. This figure is lower than our average within-person text similarity, which makes sense given that within-person changes are likely to be smaller than across-person differences. We see, moreover, that this textual similarity is stable over time, indicating that high- and low-skill workers do not appear to be trending differently in the absence of AI assistance.

Turning to the red diamonds, we see that, post-AI adoption, high- and low-skilled workers begin to use language that is more similar. The magnitude of this change—moving from 0.55 similarity to 0.61 similarity—may appear small, but given that the average within-person similarity for high-skill workers is around 0.67, this result suggests that AI assistance is associated with a substantial narrowing of language gaps.

Together, the patterns in Figures 13 and 14 suggest that low-skill workers are converging toward high-skill workers, rather than the opposite. This finding is consistent with AI models disseminating the behaviors of high-skilled workers to lower-skilled workers. In such a scenario, we would expect

low-skill workers to change their communication patterns more following AI-deployment. Top performers, meanwhile, would change less because the AI model is more likely to suggest language they already use.

6 Effects on the Experience of Work

Qualitative studies suggest that working conditions for contact center agents can be unpleasant. The repetitive nature of the job, coupled with regular exposure to challenging and emotionally charged conversations, can contribute to burnout and high turnover rates. Additionally, contact center work for US-based businesses is frequently outsourced to lower-income countries such as India and the Philippines, meaning that agents often work difficult hours and may face cultural barriers or judgements when speaking with customers.

Increases in worker productivity may not necessarily lead workers to be happier with their jobs, especially if workers feel pressured to work faster and faster. In this section, we examine the impact of generative AI on one key aspect of the workplace experience: how agents are treated by customers, as measured by customer sentiment and requests to speak with a manager. We also examine the impact of AI assistance on worker turnover as an overall indicator of worker satisfaction.

6.1 Customer Sentiment

Customers often vent their frustrations on anonymous service agents and, in our data, we see regular instances of swearing, verbal abuse, and “yelling” (typing in all caps). Service workers are called upon to absorb such customer frustrations while limiting one’s own emotional reaction (Hochschild, 2019). The stress associated with this type of emotional labor is often cited as a key cause of burnout and attrition among customer service workers (Lee, 2015).

Access to AI-assistance may impact how customers treat agents, but the direction and magnitude of these impacts are ambiguous. AI assistance may improve the tenor of conversations by helping agents set customer expectations or resolve their problems more quickly. Alternatively, customers may become more frustrated if AI-suggested language feels “corporate” or insincere.

To assess this, we attempt to capture the affective nature of both agent and customer text, using sentiment analysis (Mejova, 2009). For this analysis, we use SiEBERT, an LLM that is fine-tuned for sentiment analysis using a variety of datasets, including product reviews and tweets (Hartmann et al., 2023). Sentiment is measured on a scale from -1 to 1 , where -1 indicates negative sentiment and 1 indicates positive. In a given conversation, we compute separate sentiment scores for both

agent and customer text. We then aggregate these chat-level variables into a measure of average agent sentiment and average customer sentiment for each agent-year-month.

Panel A of Figure 15 shows the distribution of customer sentiment scores. On average, customer sentiments in our data are mildly positive and normally distributed around a mean of 0.14, except for a mass of very positive and very negative scores. Panel B shows the distribution of sentiments associated with agents: agents are unfailingly positive, with a mean sentiment score of 0.89. This reflects the fact that agents are trained to be extremely polite and friendly, even prior to AI access.

Panels C and D consider how sentiment scores respond following the roll-out of AI assistance. In Panel C, we see an immediate and persistent improvement in customer sentiment. This effect is economically large: according to Column 1 of Table 4, access to AI improves the mean customer sentiments (averaged over an agent-month) by 0.18 points, equivalent to half of a standard deviation. In Panel D, we see no detectable effect for agent sentiment, which is already very high at baseline. Column 2 of Table 4 indicates that agent sentiments increase by only 0.02 points or about 1% of a standard deviation.

Focusing on customer sentiment, Appendix Figure A.8 examines whether access to AI has different impacts for across agents. We find that access to AI assistance significantly improves how customers treat agents of all skill and experience levels, with the largest effects for agents in the lower to lower-middle range of both the skill and tenure distributions. Consistent with our productivity results, the highest-performing and most-experienced agents see the smallest benefits of AI access. These results suggest that AI recommendations, which were explicitly designed to prioritize more empathetic responses, may improve agents' demonstrated social skills and have a positive emotional impact on customers.

6.2 Customer Confidence and Managerial Escalation

Changes in individual worker-level productivity may have broader implications for organizational workflows (Garicano, 2000; Athey et al., 1994; Athey and Stern, 1998). In most customer service settings, front-line agents attempt to resolve customer problems but can seek the help of supervisors when they are unsure of how to proceed. Customers, knowing this, will sometimes attempt to escalate a conversation by asking to speak to a manager. This type of request generally occurs when the customer feels that the current agent is not equipped to address their problem or becomes frustrated.

In Figure 16, consider the impact of access to AI-assistance on the frequency of chat escalation. The outcome variable we focus on is the share of an agent's chats in which a customer requests to speak to a manager or supervisor, aggregated to the year-month level. We focus on requests

for escalation rather than actual escalations both because we lack data on actual escalations and because requests are a better measure of customer confidence in an agent’s competence or authority. Following the introduction of AI assistance, we see a gradual decline in requests for escalation. Relative to a baseline rate of approximately 6 percentage points, these coefficients suggest that AI assistance generates an almost 25% decline in customer requests to speak to a manager. In Appendix Figure A.9, we consider how these patterns change by the skill and experience of the worker. Consistent with our other results, we find that requests for escalation are disproportionately reduced for agents who were less skilled or less experienced at the time of AI adoption.

6.3 Attrition

The adoption of generative AI tools can have a variety of impacts on workers: their productivity, the amount of stress they encounter on the job, and how they are perceived by customers, to name a few. While we cannot observe all these factors, we can look at turnover patterns to provide one overarching measure of how workers are impacted by AI technology at work.

For this analysis, we compare attrition rates among treated agents to those of untreated agents with the same tenure. We drop observations for treated agents before treatment because they do not experience attrition by construction (they must survive to be treated in the future). Our analysis also controls for location and time fixed effects.

Figure 17 plots the impact of AI access on attrition: Panel A considers how this varies by agent tenure while Panel B considers heterogeneity by agent skill. Consistent with our findings so far, Panel A shows that access to AI assistance is associated with the strongest reductions in attrition among newer agents, those with less than 6 months of experience. The magnitude of this coefficient, around 10 percentage points, translates into a 40% decrease relative to a baseline attrition rate in this group of 25%. In Panel B, we examine attrition by worker skill. Here, we find a significant decrease in attrition for all skill groups, but no systematic gradient.

Finally, we note that these results should be taken with more caution relative to our main results because we are unable to include agent fixed effects to control for unobservable differences between agents with and without access to AI assistance. This is because attrition can only occur once for any given individual. Our results may overstate the impact of AI access on attrition if, for example, access to the AI tool is more likely to be given to agents whom the firm believes are more likely to stay.

7 Conclusion

Advancements in AI technologies open up a broad set of economic possibilities. Our paper provides the first empirical evidence of the effects of a generative AI tool in a real-world workplace. In our setting, we find that access to AI-generated recommendations increases worker productivity, improves customer sentiment, and is associated with reductions in employee turnover.

We hypothesize that part of the effect we document is driven by the AI system’s ability to embody the best practices of high-skill workers in our firm and make it accessible to other workers. These practices may have previously been difficult to disseminate because they involve tacit knowledge. Consistent with this, we see that AI assistance leads to substantial improvements in problem resolution and customer satisfaction for newer- and less-skilled workers but does not help the highest-skilled or most-experienced workers on these measures. Furthermore, agents who have used the system perform somewhat better even when the system is unexpectedly disabled. Analyzing the text of agent conversations, we find suggestive evidence that AI recommendations lead low-skill workers to communicate more like high-skill workers.

Our findings, and their limitations, point to a variety of directions for future research.

Most importantly, our results do not capture the potential longer-term impacts of generative AI on skill demand, job design, wages, or customer demand. It is unclear, for example, whether improvements in customer service productivity will lead to more or less demand for customer service workers. If the demand for customer support is inelastic, then generative AI tools may reduce demand and wages in this sector in the long run. Alternatively, better product support could lead customers to seek out representatives for a wider range of questions; this, in turn, could increase demand for workers or give them new responsibilities, such as collecting customer feedback for the product development team (Berg et al., 2018; Korinek, 2022).

Our findings also raise questions about the nature of worker productivity. Traditionally, a support agent’s productivity refers to their ability to help the customers they come in contact with. Yet, in a setting where customer service conversations are fed into training datasets, a worker’s productivity also includes their ability to provide ML models with examples of successful behaviors that can be shared with others. In our setting, top performers contribute many of the examples used to train the AI system we study, but they see relatively few improvements in their own productivity as a result. Under our data firm’s current pay practices, these workers may even see a reduction in their pay because bonuses are calculated relative to other agents’ performance. Our results therefore raise questions about how workers, particularly top performers, should be compensated for the data that they provide to AI systems.

Finally, as a potential general-purpose technology, generative AI can and will be deployed in a variety of ways, and the effects we find may not generalize across all firms and production processes (Eloundou et al., 2023). For example, our setting has a relatively stable product and a set of technical support questions. In areas where the product or environment is changing rapidly, the relative value of AI recommendations may be different: they may be better able to synthesize changing best practices, or they may actually impede learning by promoting outdated practices observed in historical training data.

Given the early stage of generative AI, these and other questions deserve further scrutiny.

References

- Acemoglu, Daron and David Autor**, “Skills, tasks and technologies: Implications for employment and earnings,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 1043–1171.
- **and Pascual Restrepo**, “Low-Skill and High-Skill Automation,” *Journal of Human Capital*, June 2018, 12 (2), 204–232.
- **and –**, “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 2020, 128 (6), 2188–2244. _eprint: <https://doi.org/10.1086/705716>.
- **, Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas**, “Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey,” 2022.
- **, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti**, “Technology, Information, and the Decentralization of the Firm*,” *The Quarterly Journal of Economics*, November 2007, 122 (4), 1759–1799. _eprint: <https://academic.oup.com/qje/article-pdf/122/4/1759/5234557/122-4-1759.pdf>.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad**, “The Skill Complementarity of Broadband Internet *,” *The Quarterly Journal of Economics*, July 2015, 130 (4), 1781–1824. _eprint: <https://academic.oup.com/qje/article-pdf/130/4/1781/30637431/qjv028.pdf>.
- Athey, Susan and Scott Stern**, “An Empirical Framework for Testing Theories About Complementarity in Organizational Design,” Working Paper 6600, National Bureau of Economic Research June 1998.
- **and –**, “The Impact of Information Technology on Emergency Health Care Outcomes,” *RAND Journal of Economics*, Autumn 2002, 33 (3), 399–432.
- **, Joshua Gans, Scott Schaefer, and Scott Stern**, “The Allocation of Decisions in Organizations,” *Stanford Graduate School of Business*, 1994.
- Autor, David**, “Polanyi’s Paradox and the Shape of Employment Growth,” Working Paper w20485, National Bureau of Economic Research September 2014.
- Autor, David H., Frank Levy, and Richard J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 2003, 118 (4), 1279–1333.
- **, Lawrence F. Katz, and Alan B. Krueger**, “Computing Inequality: Have Computers Changed the Labor Market?,” *The Quarterly Journal of Economics*, November 1998, 113 (4), 1169–1213. _eprint: <https://academic.oup.com/qje/article-pdf/113/4/1169/5406877/113-4-1169.pdf>.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson**, “Artificial Intelligence, Firm Growth, and Product Innovation,” May 2022.
- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio**, “Neural Machine Translation by Jointly Learning to Align and Translate,” in Yoshua Bengio and Yann LeCun, eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Baker, George P. and Thomas N. Hubbard**, “Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information,” *American Economic Review*, June 2003, 93 (3), 551–572.
- Bartel, Ann, Casey Ichniowski, and Kathryn Shaw**, “How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills*,” *The Quarterly Journal of Economics*, 11 2007, 122 (4), 1721–1758.

- Berg, Jeff, Avinash Das, Vinay Gupta, and Paul Kline**, “Smarter call-center coaching for the digital world,” Technical Report, McKinsey & Company November 2018.
- Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen**, “The Distinct Effects of Information Technology and Communication Technology on Firm Organization,” *Management Science*, 2014, 60 (12), 2859–2885.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event Study Designs: Robust and Efficient Estimation,” 2022.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt**, “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *The Quarterly Journal of Economics*, 02 2002, 117 (1), 339–376.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei**, “Language Models are Few-Shot Learners,” July 2020. arXiv:2005.14165 [cs].
- Brynjolfsson, Erik and Tom Mitchell**, “What Can Machine Learning, Do? Workforce Implications,” *Science*, December 2017, 358, 1530–1534.
- Bubeck, Sebastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg et al.**, “Sparks of artificial general intelligence: Early experiments with gpt-4,” *arXiv preprint arXiv:2303.12712*, 2023.
- Buesing, Eric, Vinay Gupta, Sarah Higgins, and Raelyn Jacobson**, “Customer care: The future talent factory,” Technical Report, McKinsey & Company June 2020.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with multiple time periods,” *Journal of Econometrics*, December 2021, 225 (2), 200–230.
- Calvino, Flavio and Luca Fontanelli**, “A Portrait of AI Adopters across Countries: Firm Characteristics, Assets’ Complementarities and Productivity,” Technical Report, OECD, Paris April 2023.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer**, “The Effect of Minimum Wages on Low-Wage Jobs*,” *The Quarterly Journal of Economics*, May 2019, 134 (3), 1405–1454.
- Choi, Jonathan H. and Daniel Schwarcz**, “AI Assistance in Legal Analysis: An Empirical Study,” August 2023.
- Chui, Michael, Bryce Hall, Alex Singla, and Alex Sukharevsky**, “Global survey: The state of AI in 2021,” Technical Report, McKinsey & Company 2021.
- de Chaisemartin, Clément and Xavier D’Haultfoeulle**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, September 2020, 110 (9), 2964–96.
- Dunn, Andrew, Diana Inkpen, and Răzvan Andonie**, “Context-Sensitive Visualization of Deep Learning Natural Language Processing Models,” 2021.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock**, “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models,” March 2023. arXiv:2303.10130 [cs, econ, q-fin].
- Felten, Edward W., Manav Raj, and Robert Seamans**, “Occupational Heterogeneity in Exposure to Generative AI,” April 2023.

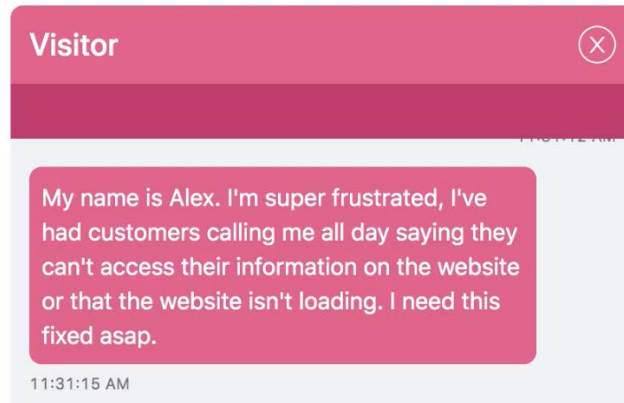
- Garicano, Luis**, “Hierarchies and the Organization of Knowledge in Production,” *Journal of Political Economy*, 2000, 108 (5), 874–904. Publisher: The University of Chicago Press.
- **and Esteban Rossi-Hansberg**, “Knowledge-Based Hierarchies: Using Organizations to Understand the Economy,” *Annual Review of Economics*, 2015, 7 (1), 1–30.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, December 2021, 225 (2), 254–277.
- Google**, “AI vs. Machine Learning: How Do They Differ?”
- Gretz, Whitney and Raelyn Jacobson**, “Boosting contact-center performance through employee engagement,” Technical Report, McKinsey & Company 2018.
- Hartmann, Jochen, Mark Heitmann, Christian Siebert, and Christina Schamp**, “More than a Feeling: Accuracy and Application of Sentiment Analysis,” *International Journal of Research in Marketing*, 2023, 40 (1), 75–87.
- Hochschild, Arlie Russell**, *The managed heart: Commercialization of human feeling*, University of California press, 2019.
- Hoffman, Mitchell, Lisa B Kahn, and Danielle Li**, “Discretion in Hiring*,” *The Quarterly Journal of Economics*, 10 2017, 133 (2), 765–800.
- Hugging Face**, “sentence-transformers/all-MiniLM-L6-v2,” April 2023.
- Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei**, “Scaling laws for neural language models,” *arXiv preprint arXiv:2001.08361*, 2020.
- Katz, Lawrence F. and Kevin M. Murphy**, “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 1992, 107 (1), 35–78.
- Korinek, Anton**, “How innovation affects labor markets: An impact assessment,” Working Paper, Brookings Institution June 2022.
- Koroteev, M. V.**, “BERT: A Review of Applications in Natural Language Processing and Understanding,” 2021.
- Lee, Don**, “The Philippines has become the call-center capital of the world,” *Los Angeles Times*, February 2015. Section: Business.
- Legg, Shane, Marcus Hutter et al.**, “A collection of definitions of intelligence,” *Frontiers in Artificial Intelligence and applications*, 2007, 157, 17.
- Li, Chun**, “OpenAI’s GPT-3 Language Model: A Technical Overview,” June 2020.
- Liu, Yiheng, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, Zihao Wu, Dajiang Zhu, Xiang Li, Ning Qiang, Dingang Shen, Tianming Liu, and Bao Ge**, “Summary of ChatGPT/GPT-4 Research and Perspective Towards the Future of Large Language Models,” April 2023. arXiv:2304.01852 [cs].
- Meijer, Erik**, “Behind every great deep learning framework is an even greater programming languages concept (keynote),” in “Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering” 2018, pp. 1–1.
- Mejova, Yelena**, “Sentiment Analysis: An Overview,” *University of Iowa, Computer Science Department*, 2009.
- Michaels, Guy, Ashwini Natraj, and John Van Reenen**, “Has ICT Polarized Skill Demand? Evidence from Eleven Countries Over Twenty-Five Years,” *The Review of Economics and Statistics*, 2014, 96 (1), 60–77.

- Nguyen, Nhan and Sarah Nadi**, “An Empirical Evaluation of GitHub Copilot’s Code Suggestions,” in “2022 IEEE/ACM 19th International Conference on Mining Software Repositories (MSR)” May 2022, pp. 1–5. ISSN: 2574-3864.
- Noy, Shakked and Whitney Zhang**, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Available at SSRN 4375283*, 2023.
- OECD**, *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*, Paris: Organisation for Economic Co-operation and Development, 2023.
- OpenAI**, “GPT-4 Technical Report,” Technical Report, OpenAI March 2023.
- Ouyang, Long, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe**, “Training language models to follow instructions with human feedback,” March 2022. arXiv:2203.02155 [cs].
- Patel, Dylan and Gerald Wong**, “GPT-4 Architecture, Infrastructure, Training Dataset, Costs, Vision, MoE,” 2023.
- Peng, Baolin, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao**, “Check Your Facts and Try Again: Improving Large Language Models with External Knowledge and Automated Feedback,” 2023.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer**, “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” 2023.
- Polanyi, Michael**, *The Tacit Dimension*, Chicago, IL: University of Chicago Press, May 1966.
- Radford, Alec and Karthik Narasimhan**, “Improving Language Understanding by Generative Pre-Training,” 2018.
- , **Jeff Wu, Rewon Child, D. Luan, Dario Amodei, and Ilya Sutskever**, “Language Models are Unsupervised Multitask Learners,” 2019.
- Roose, Kevin**, “A Conversation With Bing’s Chatbot Left Me Deeply Unsettled,” *The New York Times*, February 2023.
- Rosen, Sherwin**, “The Economics of Superstars,” *The American Economic Review*, 1981, 71 (5), 845–858.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- Syverson, Chad**, “What Determines Productivity?,” *Journal of Economic Literature*, June 2011, 49 (2), 326–65.
- Taniguchi, Hiroya and Ken Yamada**, “ICT Capital-Skill Complementarity and Wage Inequality: Evidence from OECD Countries,” *Labour Economics*, June 2022, 76, 102151. arXiv:1904.09857 [econ, q-fin].
- The White House**, “The Impact of Artificial Intelligence on the Future of Workforces in the European Union and the United States of America,” Technical Report, The White House December 2022.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin**, “Attention Is All You Need,” December 2017. arXiv:1706.03762 [cs].
- Zhao, Shuyin**, “GitHub Copilot now has a better AI model and new capabilities,” February 2023.

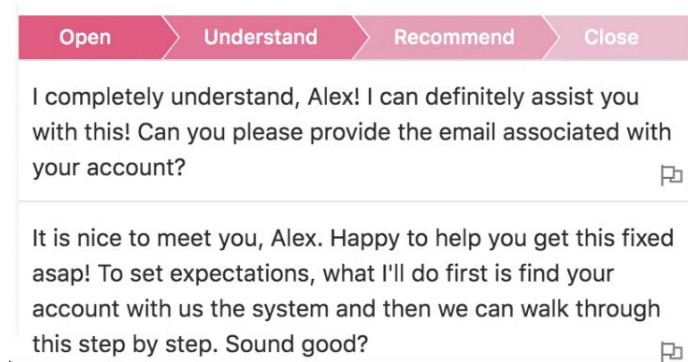
Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David Beede, Catherine Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz, “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey,” Working Papers 20-40, Center for Economic Studies, U.S. Census Bureau December 2020.

FIGURE 1: SAMPLE AI OUTPUT

A. SAMPLE CUSTOMER ISSUE




B. SAMPLE AI-GENERATED SUGGESTED RESPONSE




NOTES: This figure shows sample suggestions of output generated by the AI model. The suggested responses are only visible to the agent. Workers can choose to ignore, accept or somewhat incorporate the AI suggestions into their response to the customer.

FIGURE 2: RAW PRODUCTIVITY DISTRIBUTIONS, BY AI TREATMENT




NOTES: This figure shows the distribution various outcome measures. We split this sample into agent-month observations for agents who eventually receive access to the AI system before deployment (“Pre AI”), after deployment (“Post AI”), and for agent-months associated with agents who never receive access (“Never AI”). Our primary productivity measure is “resolutions per hour,” the number of customer issues the agent is able to successfully resolve per hour. We also provide descriptives for “average handle time,” the average length of time an agent takes to finish a chat; “chats per hour,” the number of chats completed per hour incorporating multitasking; “resolution rate,” the share of conversations that the agent is able to resolve successfully; and “net promoter score” (NPS), which are calculated by randomly surveying customers after a chat and calculating the percentage of customers who would recommend an agent minus the percentage who would not. All data come from the firm’s software systems.

FIGURE 3: EVENT STUDIES, RESOLUTIONS PER HOUR



NOTES: These figures plot the coefficients and 95% confidence intervals from event study regressions of AI model deployment using the [Sun and Abraham \(2021\)](#) interaction weighted estimator. See text for additional details. Panel A plots the resolutions per hour and Panel B plots the natural log of the measure. All specifications include agent and chat year-month, location, agent tenure and company fixed effects. Robust standard errors are clustered at the agent level.


FIGURE 4: EVENT STUDIES, ADDITIONAL OUTCOMES



NOTES: These figures plot the coefficients and 95% confidence intervals from event study regressions of AI model deployment using the [Sun and Abraham \(2021\)](#) interaction weighted estimator. See text for additional details. Panel A plots the average handle time or the average duration of each technical support chat. Panel B plots the number of chats an agent completes per hour, incorporating multitasking. Panel C plots the resolution rate, the share of chats successfully resolved, and Panel D plots net promoter score, which is an average of surveyed customer satisfaction. All specifications include agent and chat year-month, location, agent tenure and company fixed effects. Robust standard errors are clustered at the agent level.

FIGURE 5: HETEROGENEITY OF AI IMPACT, BY SKILL AND TENURE

A. IMPACT OF AI ON RESOLUTIONS PER HOUR, BY SKILL AT DEPLOYMENT




B. IMPACT OF AI ON RESOLUTIONS PER HOUR, BY TENURE AT DEPLOYMENT




NOTES: These figures plot the impacts of AI model deployment on log(resolutions per hour) for different groups of agents. Agent skill is calculated as the agent's trailing three month average of performance on average handle time, call resolution, and customer satisfaction, the three metrics our firm uses to assess agent performance. Within each month and company, agents are grouped into quintiles, with the most productive agents in quintile 5 and the least productive in quintile 1. Pre-AI worker tenure is the number of months an agent has been employed when they receive access to AI recommendations. All specifications include agent and chat year-month, location, and company fixed effects and standard errors are clustered at the agent level. Panel A includes controls for agent tenure at deployment and Panel B includes controls for agent skill at deployment.

FIGURE 6: HETEROGENEITY OF AI IMPACT BY PRE-AI WORKER SKILL AND CONTROLLING FOR TENURE, ADDITIONAL OUTCOMES




NOTES: These figures plot the impacts of AI model deployment on four measures of productivity and performance, by pre-deployment worker skill. Agent skill is calculated as the agent's trailing three month average of performance on average handle time, call resolution, and customer satisfaction, the three metrics our firm uses for agent performance. Within each month and company, agents are grouped into quintiles, with the most productive agents within each firm in quintile 5 and the least productive in quintile 1. Panel A plots the average handle time or the average duration of each technical support chat. Panel B graphs chats per hour, or the number of chats an agent can handle per hour. Panel C plots the resolution rate, and Panel D plots net promoter score, an average of surveyed customer satisfaction. All specifications include agent and chat year-month, location, and company fixed effects, controls for agent tenure and standard errors are clustered at the agent level.

FIGURE 7: HETEROGENEITY OF AI IMPACT BY PRE-AI WORKER TENURE CONTROLLING FOR SKILL, ADDITIONAL OUTCOMES




NOTES: These figures plot the impacts of AI model deployment on measures of productivity and performance by pre-AI worker tenure, defined as the number of months an agent has been employed when they receive access to the AI model. Panel A plots the average handle time or the average duration of each technical support chat. Panel B graphs chats per hour, or the number of chats an agent can handle per hour. Panel C plots the resolution rate, and Panel D plots net promoter score, an average of surveyed customer satisfaction. All specifications include agent and chat year-month, location, and company fixed effects, controls for agent skill at deployment and standard errors are clustered at the agent level.

FIGURE 8: EXPERIENCE CURVES BY DEPLOYMENT COHORT



NOTES: This figure plot the relationship between productivity and job tenure. The red line plots the performance of always-treated agents, those who have access to AI assistance from their first month on the job. The blue line plots agents who are never treated. The green line plots agents who spend their first four months of work without the AI assistance, and gain access to the AI model during their fifth month on the job. 95% confidence intervals are shown.


FIGURE 9: HETEROGENEITY OF AI IMPACT, BY AI ADHERENCE




NOTES: Panel A plots the distribution of AI adherence, averaged at the agent-month level, weighted by the log of the number of AI recommendations for that agent-month. Panel B shows the impact of AI assistance on resolutions by hour, by agents grouped by their initial adherence, defined as the share of AI recommendations they followed in the first month of treatment.

FIGURE 10: AI ADHERENCE OVER TIME


A. BY ADHERENCE AT AI MODEL DEPLOYMENT



B. BY AGENT TENURE AT AI MODEL DEPLOYMENT




C. BY AGENT SKILL AT AI MODEL DEPLOYMENT




NOTES: This figure plots the share of AI suggestions followed by agents as a function of the number of months each agent has had access to the AI model. In Panel A, we divide agents into terciles based on their adherence to AI suggestions in the first month. In Panel B, we divide agents into groups based on their tenure at the firm at the time of AI model deployment. In Panel C, we divide workers into terciles of pre-deployment productivity as defined by our skill index. All data come from the firm's internal software systems.

FIGURE 11: PRODUCTIVITY DURING AI SYSTEM OUTAGES



NOTES: This figure plots event studies of the impact of the roll out of AI-assistance on chat times at the individual chat level. Panel A restricts to post-treatment chats that do not occur during any period where there is a system outage. Panel B restricts to post-treatment chats that only occur during a large system outage. Standard errors are clustered at the agent level.


FIGURE 12: PRODUCTIVITY DURING AI SYSTEM OUTAGES, BY INITIAL AI ADHERENCE




NOTES: This figure plots event studies for the impact of AI system rollout of chat duration. All post-adoption data is restricted to periods in which there is a major outage. Panel A restricts to only chats generated by ever-treated agents who with high initial AI adherence (top tercile) while Panel B restricts to agents with low initial adherence (bottom tercile). Agents who are never treated are excluded from this analysis.

FIGURE 13: WITHIN AGENT TEXTUAL ANALYSIS

A. WITHIN-PERSON TEXTUAL SIMILARITY TO MONTH PRIOR TO AI




B. WITHIN-PERSON TEXTUAL CHANGE, LOW VS. HIGH SKILL




NOTES: Panel A plots the average similarity between an agent's chats each week and a comparison group of their conversations in the month prior to AI deployment. To avoid comparing conversations to themselves, we exclude calculating the similarity in the month prior to deployment. Panel B plots the average difference between an agent's pre-AI corpus of chat messages and that same agent's post-AI corpus, controlling for year-month and agent tenure. The first bar represents the average pre-post text difference for agents in the highest quintile of pre-AI skill, as measured by a weighted index of their chats per hour, resolution rate, and customer satisfaction score. The low-skill bar represents the same type of pre-post text difference among the lowest skill quintile. Agent skill, or relative productivity, is defined at the time of treatment.

FIGURE 14: TEXT SIMILARITY BETWEEN LOW-SKILL AND HIGH-SKILL WORKERS, PRE AND POST AI




NOTES: This figure plots the average text similarity between the top and bottom quintile of agents. The blue line plots the similarity for never treated or pre-treatment agents, the red line plots the similarity for agents with access to the AI model. For agents in the treatment group, we define agent skill prior to AI model deployment. Our analysis includes controls for agent tenure.

FIGURE 15: CONVERSATION SENTIMENT




NOTES: Each panel of this figure plots the impact of AI model deployment on conversational sentiment. Panel A shows average customer sentiments. Panel B shows average agent sentiments. Panel C plots the event study of AI model deployment on customer sentiment and Panel D plots the corresponding estimate for agent sentiment. Sentiment is measured using SiEBERT, a fine-tuned checkpoint of a RoBERTA, an English language transformer model. All data come from the firm's internal software systems.

FIGURE 16: IMPACT OF AI ON CHAT ESCALATION



NOTES: This figure reports the coefficients and 95% confidence intervals for the event study of AI model deployment on requests for manager assistance. Standard errors are clustered at the agent level.

FIGURE 17: IMPACT OF AI MODEL DEPLOYMENT ON WORKER ATTRITION



NOTES: This figure presents the results of the impact of AI model deployment on workers' likelihood of attrition. Panel A graphs the effects of AI assistance on attrition by agent tenure at AI model deployment. Panel B plots the same impact by agent skill index at AI model deployment. All specifications include chat year and month fixed effects, as well as agent location, company and agent tenure. All robust standard errors are clustered at the agent level. All data come from the firm's internal software systems.

TABLE 1: APPLICANT SUMMARY STATISTICS

Variable	All	Never Treated	Treated, Pre	Treated, Post
Chats	3,007,501	945,954	882,105	1,180,446
Agents	5,179	3,523	1,341	1,636
Number of Teams	133	111	80	81
Share US Agents	.11	.15	.081	.072
Distinct Locations	25	25	18	17
Average Chats per Month	127	83	147	188
Average Handle Time (Min)	41	43	43	35
St. Average Handle Time (Min)	23	24	24	22
Resolution Rate	.82	.78	.82	.84
Resolutions Per Hour	2.1	1.7	2	2.5
Customer Satisfaction (NPS)	79	78	80	80

NOTES: This table shows conversations, agent characteristics and issue resolution rates, customer satisfaction and average call duration. The sample in Column 1 consists of all agents in our sample. Column 2 includes control agents who were never given access to the AI model. Column 3 and 4 present pre-and-post AI model deployment summary statistics for treated agents who were given access to the AI model. All data come from the firm's internal software systems.

TABLE 2: MAIN EFFECTS: PRODUCTIVITY (RESOLUTIONS PER HOUR)

VARIABLES	(1) Res./Hr	(2) Res./Hr	(3) Res./Hr	(4) Log(Res./Hr)	(5) Log(Res./Hr)	(6) Log(Res./Hr)
Post AI X Ever Treated	0.468*** (0.0542)	0.371*** (0.0520)	0.301*** (0.0498)	0.221*** (0.0211)	0.180*** (0.0188)	0.138*** (0.0199)
Ever Treated	0.109* (0.0582)			0.0572* (0.0316)		
Observations	13,225	12,328	12,328	12,776	11,904	11,904
R-squared	0.250	0.563	0.575	0.260	0.571	0.592
Year Month FE	YES	YES	YES	YES	YES	YES
Location FE	YES	YES	YES	YES	YES	YES
Agent FE	-	YES	YES	-	YES	YES
Agent Tenure FE	-	-	YES	-	-	YES
DV Mean	2.121	2.174	2.174			

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.10

NOTES: This table presents the results of difference-in-difference regressions estimating the impact of AI model deployment on our main measure of productivity, resolutions per hour, the number of technical support problems resolved by an agent per hour (res/hour). Columns 1 and 4 include agent geographic location and year-by-month fixed effects. Columns 2 and 5 include agent-level fixed effects, and columns 3 and 6, our preferred specification, also control for agent tenure. All standard errors are clustered at the agent location level. All data come from the firm's internal software systems.

TABLE 3: MAIN EFFECTS: ADDITIONAL OUTCOMES

VARIABLES	(1) AHT	(2) Calls/Hr	(3) Res. Rate	(4) NPS	(5) Log(AHT)	(6) Log(Calls/Hr)	(7) Log(Res. Rate)	(8) Log(NPS)
Post AI X Ever Treated	-3.750*** (0.476)	0.366*** (0.0363)	0.0128* (0.00717)	-0.128 (0.660)	-0.0851*** (0.0110)	0.149*** (0.0142)	0.00973* (0.00529)	-0.000406 (0.00915)
Observations	21,885	21,885	12,328	12,578	21,885	21,885	11,904	12,188
R-squared	0.590	0.564	0.369	0.525	0.622	0.610	0.394	0.565
Year Month FE	YES	YES	YES	YES	YES	YES	YES	YES
Location FE	YES	YES	YES	YES	YES	YES	YES	YES
Agent FE	YES	YES	YES	YES	YES	YES	YES	YES
Agent Tenure FE	YES	YES	YES	YES	YES	YES	YES	YES
DV Mean	40.65	2.557	0.821	79.58				

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.10

NOTES: This table presents the results of difference-in-difference regressions estimating the impact of AI model deployment on measures of productivity and agent performance. Post AI X Treated measures the impact of AI model deployment after deployment on treated agents for average handle time or average call duration, chats per hour, the number of chats an agent handles per hour, resolution rate, the share of technical support problems they can resolve and net promoter score (NPS), an estimate of customer satisfaction, and each metrics corresponding natural log equivalents. All specifications include agent fixed effects and chat year and month fixed effects, as well as controls for agent location and agent tenure. All standard errors are clustered at the agent location level. All data come from the firm's internal software systems.

TABLE 4: AGENT AND CUSTOMER SENTIMENT

VARIABLES	(1) Mean(Customer Sentiment)	(2) Mean(Agent Sentiment)
Post AI X Ever Treated	0.177*** (0.0133)	0.0198*** (0.00315)
Observations	21,218	21,218
R-squared	0.485	0.596
Year Month FE	YES	YES
Location FE	YES	YES
Agent FE	YES	YES
Agent Tenure FE	YES	YES
DV Mean	0.141	0.896

Robust standard errors in parentheses

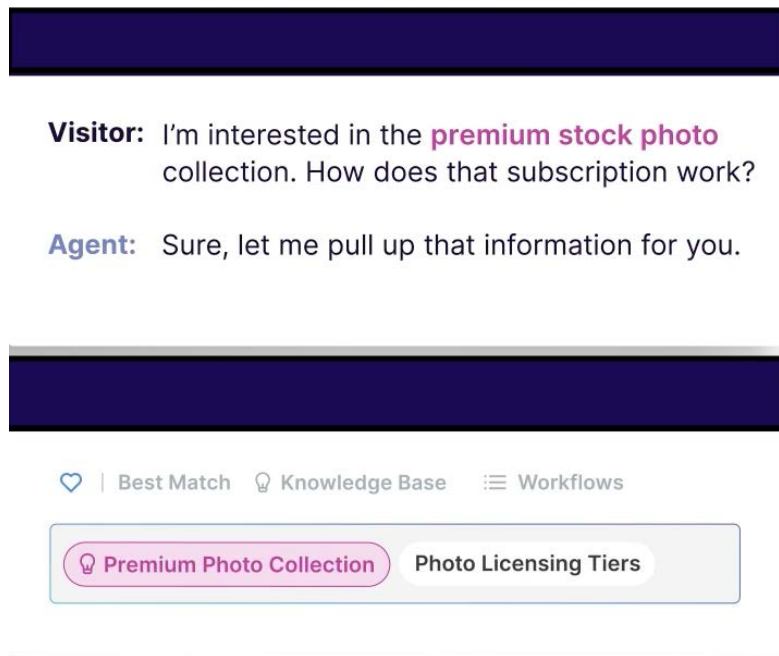
*** p<0.01, ** p<0.05, * p<0.10

NOTES: This table presents the results of difference-in-difference regressions estimating the impact of AI model deployment on measures of conversation sentiment. All specifications include agent fixed effects and chat year and month fixed effects, as well as agent location and agent tenure, which account for differing likelihood of attrition by agent tenure. All standard errors are clustered at the agent location level. All data come from the firm's internal software systems.

Appendix Materials


FIGURE A.1: SAMPLE AI TECHNICAL SUGGESTION

A. SAMPLE AI-GENERATED TECHNICAL LINK




NOTES: This figure shows a sample technical documentation suggestions made the by AI. Our data firm has an extensive set of documentation for their technical support agents, known as the knowledge base, which is like an internal company Wikipedia for product and process information. The AI will attempt to surface the most helpful technical documentation page when triggered to do so during a customer interaction. These links are only visible to the agent and agents must review to see if the resource is helpful. Workers can choose to read the suggested technical documentation or ignore the recommendation.

FIGURE A.2: DEPLOYMENT TIMELINE




NOTES: This figure shows the share of agents deployed onto the AI system over the study period. Agents are deployed onto the AI system after a training session. The firm ran a small randomized control trial in August and September of 2020. All data are from the firm's internal software systems.

FIGURE A.3: EVENT STUDIES, RESOLUTIONS PER HOUR




NOTES: This table presents the effect of AI model deployment on our main productivity outcome, resolutions per hour, using a variety of robust dynamic difference-in-differences estimators introduced in [Borusyak et al. \(2022\)](#), [Callaway and Sant'Anna \(2021\)](#), [de Chaisemartin and D'Haultfoeulle \(2020\)](#) and [Sun and Abraham \(2021\)](#) and a standard two-way fixed effects regression model. All regressions include agent level, chat-year fixed effects and controls for agent tenure. Standard errors are clustered at the agent level. Because of the number of post-treatment periods and high turnover of agents in our sample, we can only estimate five months of preperiod using [Borusyak et al. \(2022\)](#) and [de Chaisemartin and D'Haultfoeulle \(2020\)](#).

FIGURE A.4: EXPERIENCE CURVES BY DEPLOYMENT COHORT, ADDITIONAL OUTCOMES



NOTES: These figures plot the experience curves of three groups of agents over their tenure, the x-axis, against five measures of productivity and performance. The red lines plot the performance of always-treated agents, those who start work in their first month with the AI and always have access to the AI suggestions. The blue line plots agents who are never treated. The green line plots agents who spend their first four months of work without the AI model, and gain access to the AI during their fifth month on the job. All panels include 95% confidence intervals.


FIGURE A.5: HETEROGENEITY OF AI IMPACT BY INITIAL AI ADHERENCE, ADDITIONAL OUTCOMES




NOTES: These figures plot the impact of AI model deployment on additional measures of performance by quintile of initial adherence, the share of AI recommendations followed in the first month of treatment. Panel A plots the average handle time or the average duration of each technical support chat. Panel B graphs chats per hour, or the number of chats an agent can handle per hour (including working on multiple chats simultaneously). Panel C plots the resolution rate, the share of chats successfully resolved, and Panel D plots NPS, or net promoter score, is an average of surveyed customer satisfaction. All specifications include agent and chat year-month, location, and company fixed effects and controls for agent tenure. All data come from the firm's internal software systems.

FIGURE A.6: WITHIN-AGENT AI ADHERENCE OVER TIME


A. BY ADHERENCE AT AI MODEL DEPLOYMENT



B. BY AGENT TENURE AT AI MODEL DEPLOYMENT




C. BY AGENT SKILL AT AI MODEL DEPLOYMENT




NOTES: This figure plots the residualized percentage of AI suggestions followed by agents as a function of the number of months each agent has had access to the AI model, after controlling for agent level fixed effects. In Panel A, we divide agents into terciles based on their adherence to AI suggestions in the first month. In Panel B, we divide agents into groups based on their tenure at the firm at the time of AI model deployment. In Panel C, we divide workers into terciles of pre-deployment productivity as defined by our skill index. All data come from the firm's internal software systems.

FIGURE A.7: SAMPLE AI OUTAGE




NOTES: This figure plots the share of post-treatment chats with no AI suggestions during a period of a documented software outage.

FIGURE A.8: HETEROGENEITY IN CUSTOMER SENTIMENT



NOTES: Each panel of this figure plots the impact of AI model deployment on the mean sentiment per conversation. Sentiment refers to the emotion or attitude expressed in the text of the customer chat and ranges from -1 to 1 where -1 indicates very negative sentiment and 1 indicates very positive sentiment. Panel A plots the effects of AI model deployment on customer sentiment by agent tenure when AI deployed and Panel B plots the impacts by agent ex-ante productivity. All data come from the firm's internal software systems. Average sentiment is measured using SiEBERT, a fine-tuned checkpoint of a RoBERTa, an english language transformer model.

FIGURE A.9: ESCALATION, HETEROGENEITY BY WORKER TENURE AND SKILL



NOTES: Panels A and B show the effects of AI on customer requests for manager assistance, by pre-AI agent skill and in by pre-AI agent tenure. All robust standard errors are clustered at the agent location level. All data come from the firm's internal software systems.

TABLE A.1: MAIN EFFECTS: PRODUCTIVITY (LOG(RESOLUTIONS PER HOUR)), ALTERNATIVE DIFFERENCE-IN-DIFFERENCE ESTIMATORS

	Point Estimate	Standard Error	Lower Bound 95% Confidence Interval	Upper Bound 95% Confidence Interval
TWFE-OLS	0.137	0.014	0.108	0.165
Borusyak-Jaravel-Spiess	0.257	0.028	0.203	0.311
Callaway-Sant’Anna	0.239	0.025	0.189	0.289
DeChaisemartin-D’Haultfoeuille	0.116	0.021	0.075	0.156
Sun-Abraham	0.237	0.037	0.165	0.308

NOTES: This table shows the impact of AI model deployment on the log of our main productivity outcome, resolutions per hour, using robust difference-in-differences estimators introduced in [Borusyak et al. \(2022\)](#), [Callaway and Sant’Anna \(2021\)](#), [de Chaisemartin and D’Haultfoeuille \(2020\)](#) and [Sun and Abraham \(2021\)](#). All regressions include agent level, chat-year fixed effects and controls for agent tenure. The standard errors are clustered at the agent level.