workato with UserEvidence

Behind the AI Hype

The 2024 State of LLMs in Business Processes

workato

Contents

Overview

Market sentiment about generative AI has entered a new era. **Early hype is fading.** Business leaders have a clearer view of its strengths and challenges.

Major press, such as the Wall Street Journal, The Economist, and others have wondered if the AI hype over-promised and under-delivered. But very few have considered a key value driver for companies: **Business processes and operations.** As legendary business thinker Tom Davenport wrote in HBR: "Most AI applications to date seek to improve a given task. But this is missing the larger picture; smart companies are viewing the introduction of AI as the rationale for a new look at end-to-end processes."

Several firms, such as McDonalds, Air Canada, and IBM have experienced high profile failures with early generative AI projects. Conventional wisdom says gen AI issues are ***model challenges***: Ways that large language models (LLMs) need improvement. However, McKinsey notes that it is "important to recognize that the model itself makes up only 15 percent of the success [of an AI project]." There is much more to an AI strategy than training data, model weights, and prompt engineering.

To create consistent business value with AI, companies must also address ***process challenges***. These include: How gen AI is activated in business processes. How AI accesses the latest data. How AI risks are governed, and how human oversight & accountability is managed.

With independent market research from UserEvidence, we set out to understand the impact of generative AI on business processes and operations. We surveyed 1,000+ business leaders, directors and above. Participants include every department in companies of every size.

The survey shows that, despite changing market sentiment:

Business leaders believe the hype:

82%

Believe genAI is the next wave of transformation in their industry

The real results are in business processes:

89%

Report real, measurable results from genAI in business processes

AI momentum is immune to the vibe shift:

Competitors are the #1 driver

Market sentiment is only one of several drivers of AI adoption

Top Findings

Business processes = big value with AI.

Over 35% have achieved revenue growth, business agility, or better decision making.

Early adopters are winning.
Early adopters are nearly 2x more likely to report increased agility, enhanced decision making, and innovation, and 4x less likely to report no measurable results.

APIs & RAG drive better results.
LLM approaches like retrieval-augmented generation (RAG) and AI activation approaches like API-first approaches have better performance.

Governance and human oversight matter. A lot.

Governance is the #1 requirement to improve trust in AI, and 95% say human oversight is important.

More AI = more demand for human intelligence.
Business leaders believe expert reviewers - rather than front-line or entry-level employees - should sign off on AI-generated content & decisions - meaning artificial intelligence increases demand for human intelligence.

APIs are the top approach to activate AI in processes.

APIs are the #1 preference to activate AI in processes in the future.

Gen AI SaaS features are not a long term AI strategy.
Just over 13% of business leaders believe SaaS AI features are their best AI strategy moving forward.

workato

CHAPTER 1:

What's Driving AI Adoption

KEY TAKEAWAYS

- Many factors drive AI
- Pressure on departments is uneven
- IT is under the most pressure

Business leaders hear the call to use generative AI in their processes loud and clear - **71% feel pressure to implement generative AI in their processes**. The most common sources of that pressure come from outside the company. **The number one source of pressure is the competition.** (Figure 1.1).

Competitors are the top driver of AI adoption.

While - as mentioned in the overview - market sentiment on generative AI is evolving, business leaders are also responding to *other* sources of pressure, including competitors, external peers, analysts, internal peers, and board members.

FIGURE 1.1: WHAT DO YOU BELIEVE IS A SOURCE OF PRESSURE (IF ANY) TO USE GEN AI IN YOUR PROCESSES (SELECT ALL THAT APPLY)

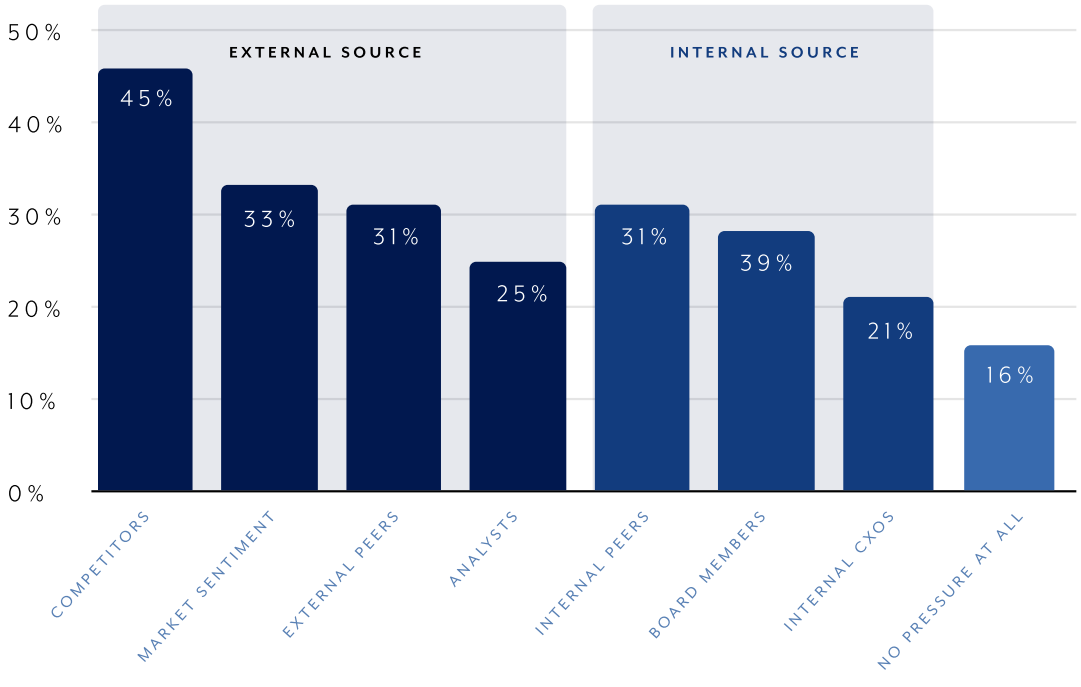

"What happens when I have human-level intelligence, for pennies on demand, accessible 24 hours, 7 days a week, 365 days a year? How does that change all the things that we do? And then the next leap from that is, how does it change what our competitors can do? That's when the real fire gets started. Because you start saying, 'Gee, if I can do this, you know, the folks that I meet in the marketplace every day, who are going for those same customers, can also do this.'"

TED SHELTON, COO, INFLECTION AI (PREVIOUSLY THE AI EXPERT PARTNER AT BAIN & COMPANY)

The pressure is not felt evenly across departments. It is being felt more acutely by IT, Finance, and HR, and less so by Sales, Marketing, and Operations (Figure 1.2).

IT, Finance, and HR are under more pressure.

FIGURE 1.2: WHICH WORD DESCRIBES THE PRESSURE YOU FEEL TO USE GENERATIVE AI IN YOUR PROCESSES? (COHORTS FROM FIGURE 0.1)

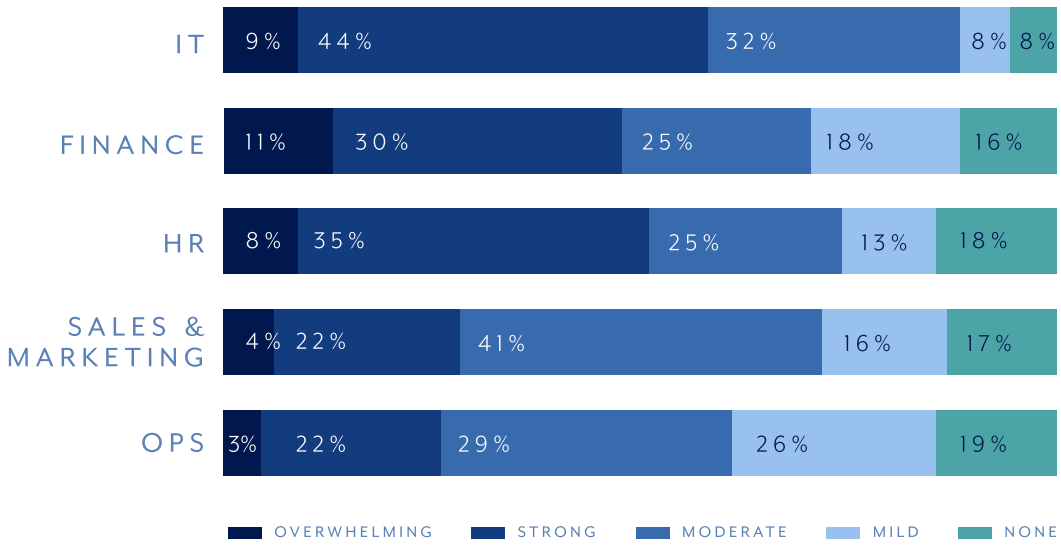

As business leaders face increasing pressure to adopt AI due to competitive forces, market expectations, and more, they are taking action. The following chapter explores the methods leaders are employing to activate AI in their operations, to transform these pressures into outcomes.

workato

CHAPTER 2:

Activating LLMs in Business Processes

KEY TAKEAWAYS

- Foundational LLMs have an early lead
- SaaS features have a bleak forecast
- APIs are the top approach

"Activating generative AI in a business process" or "AI-powered processes" means repeated, consistent use of AI in business operations such as order to cash, customer onboarding, lead generation, and more. It can take many forms: A prompt and output step in a workflow, a process reimagined by AI, or even agents that execute entire processes autonomously.

In each case, business leaders have dozens of choices to make. These include:

- **What** approach they will take to their large language models (LLMs), such as building from scratch, buying, or partnering with a vendor.
- **How** they will activate AI in the process, such as manual prompt entry, an API call, or an agent.
- **Where** in the organization they will apply it.

As we'll learn in chapter 3, these choices also shape the outcomes companies experience.

What are the top approaches to LLMs

With the popularity of OpenAI's ChatGPT, it is no surprise that business leaders are relying on industry giants. Over two-thirds (68%) said that their current approach is a foundational LLM (Figure 2.1) like OpenAI GPT, Google Gemini, or Anthropic Claude.

FIGURE 2.1: WHICH APPROACH ARE YOU TAKING TO LLMS IN YOUR PROCESSES? (SELECT ALL THAT APPLY)

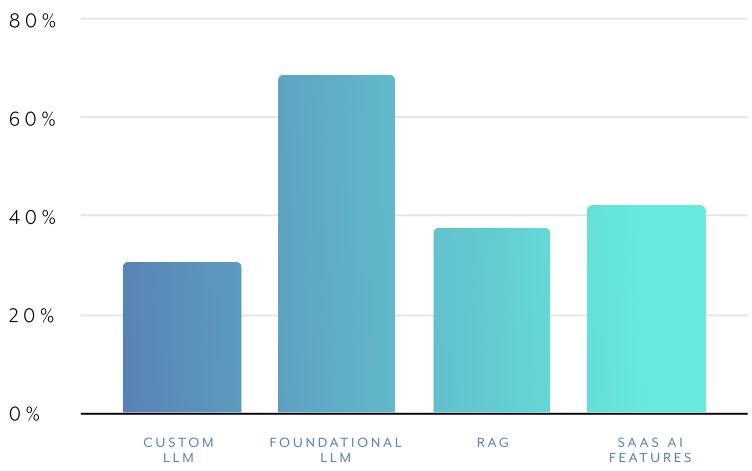

While nearly a third are building their own LLM, **a majority (67%) are using foundational models** with no modifications. **38% are customizing a foundational model** with retrieval-augmented generation (RAG).

"Gen AI is so successful compared to previous AI approaches because it's easy to start with. Foundational LLMs are good enough for many things, making them easy to adopt and appealing."

MARKUS ZIRN, CHIEF STRATEGY OFFICER, WORKATO

Nearly every SaaS vendor became an AI vendor overnight by adding generative AI features to their products. **42% of business leaders are using a SaaS app** with an AI feature (Figure 2.1). However - when asked about the future, business leaders' opinion differs from their current strategy. **Business leaders do not view SaaS features as their long term AI strategy**, with just over 13% selecting this option as their best path forward (Figure 2.2).

Only 13% see SaaS features as their long-term AI strategy.

68% of companies (Figure 2.1) are using foundational models today, which are trained on the entire internet. However, this number **drops by half when asked if it is the best approach in the future** (Figure 2.2)**.** This likely reflects a belief that models custom trained on domain expertise (such as customer support agent conversations, or IT helpdesk tickets) will be more effective. 35% of business leaders believe customized and specialized LLMs will offer the best path forward.

FIGURE 2.2: WHICH STRATEGY DO YOU BELIEVE WILL BE MOST EFFECTIVE FOR THE FUTURE? (SELECT ONE)

While retrieval-augmented generation (RAG) is used in over 38% of use cases today, **only 8% believe it is the best strategy for the future** (Figure 2.3). This is unusual. RAG is used to inform AI output with the most up-to-date context. Figure 3.3 later in this report shows RAG drives better outcomes, such as enhanced decision making. It is likely that participants believe a custom or specialized LLM in the future will be able to harness the latest context without the need for RAG. Or, when forced to choose between the two, they value a customized LLM more highly even if they still plan to use RAG in the future.

How will AI be activated

There are many ways to activate generative AI in a business process. Examples might include an employee manually prompting a model such as ChatGPT, custom code interacting with the model via API, or a robotic process automation bot pasting an output into a UI text field.

48%
believe APIs are best

13%
believe SaaS features are best

5.7%
believe RPA is best

The data shows that **the top preference of business leaders is for API-based activation of generative AI**, with over 48% saying they believe the best way to activate AI is by API-driven methods, such as custom code, an integration platform, or a specialized AI orchestration tool (Figure 2.3). Later in this report, we will see that this preference is correlated with higher performance (Figure 3.4).

FIGURE 2.3: IN THE FUTURE, WHICH METHOD TO USE GENERATIVE AI IN BUSINESS PROCESSES DO YOU BELIEVE WILL BE BEST FOR YOUR ORGANIZATION?

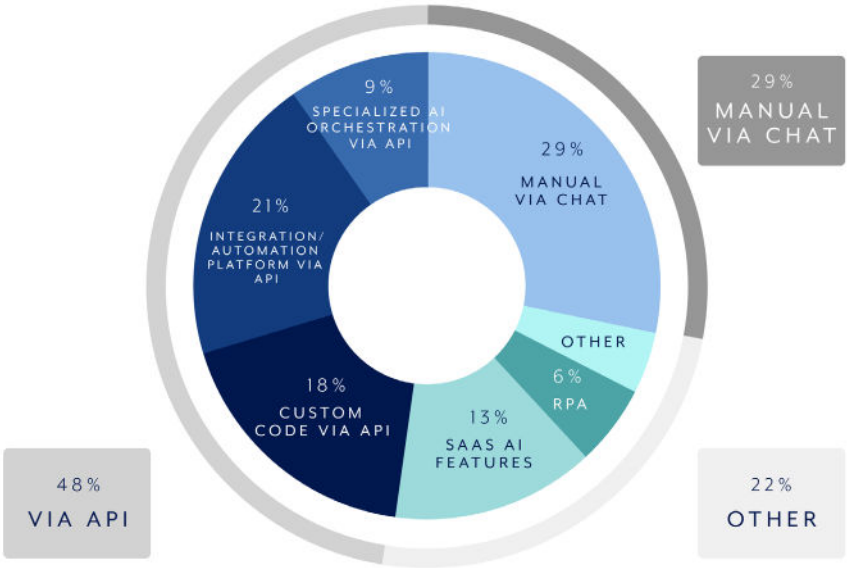

While many selected manual chat, this can either mean interacting directly with an LLM such as ChatGPT, or it can mean interacting with technology via natural language.

Notably, the AI forecast for robotic process automation (RPA) for AI-powered processes is bleak, with only 5.7% selecting it as the best approach for activating AI in the future (Figure 2.3).

Where is gen AI being used in business processes

"You can imagine dozens of use cases in HR, sales, productivity, supply chain. The list goes on and on. LLMs can make a huge difference in efficiency and productivity."

RAMA AKKIRAJU, VP OF AI AND ML FOR IT, NVIDIA

Processes in IT, customer service, operations, sales and marketing are the top early adopters of generative AI. Within the next six months, the priorities of IT and customer service shrink slightly, and HR and Finance grow, although the overall ranking remains largely the same. The pressure on HR and Finance noted in Figure 1.2 may provide some indication as to why.

FIGURE 2.4: IN WHICH DEPARTMENTS IS GENERATIVE AI BEING USED IN PROCESSES, AND WHICH ARE PLANNED WITHIN 6 MONTHS (SELECT ALL THAT APPLY)

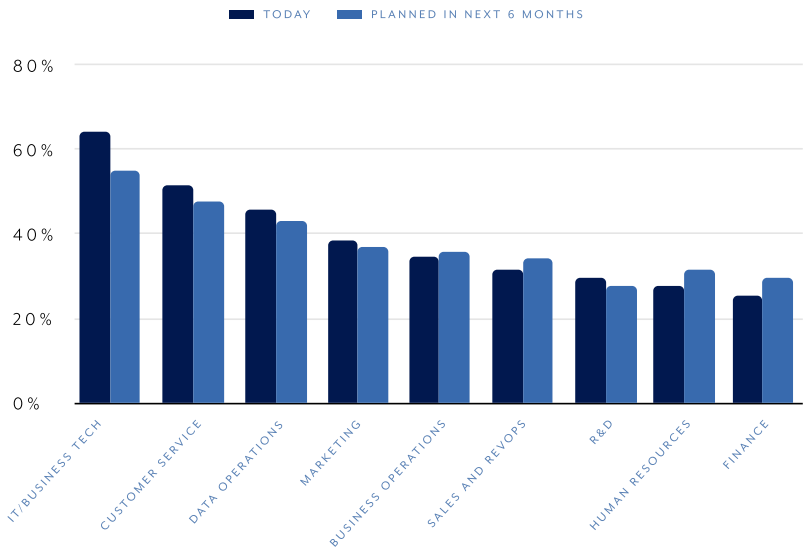

Consistently (and not surprisingly) the data shows IT as the most enthusiastic early adopters of AI in business processes. CIOs are under extraordinary pressure to deliver on their organization's AI strategy. It is also possible that IT processes represent the lowest hanging fruit, and CS, sales, and marketing bring more urgency since they are directly tied to revenue. On the other hand, Finance and HR processes may be more complex, or take longer to deliver results, leading business leaders to postpone them (Figure 2.4).

Only 23% consider AI to be fully deployed in a business process

However, the maturity of the AI implementations varies. **Less than a quarter of participants (23%) said AI is fully deployed in a business process** (Figure 2.5). As we'll see later in the report, the level of maturity has an extraordinary impact on the outcomes (Figure 3.5).

FIGURE 2.5: IF GENERATIVE AI IS BEING USED IN YOUR BUSINESS PROCESSES, TO WHAT DEGREE WOULD YOU DESCRIBE ITS USE?

The maturity of deployment varies on a department level. IT and HR appear to have made the most progress. Finance and Sales & Marketing have the least deployed and the most experimentation (Figure 2.6).

"I haven't seen many organizations put together an enterprise-wide AI strategy. Most of them are still experimenting with AI tactically or opportunistically."

MASSIMO PEZZINI, HEAD OF RESEARCH, FUTURE OF THE ENTERPRISE, WORKATO

FIGURE 2.6: IF GENERATIVE AI IS BEING USED IN YOUR BUSINESS PROCESSES, TO WHAT DEGREE WOULD YOU DESCRIBE ITS USE? (COHORTS FROM FIGURE 0.1)

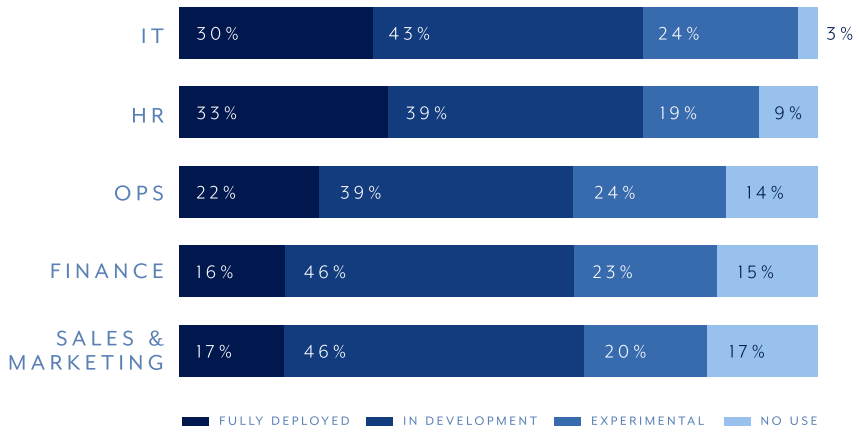

KEY TAKEAWAYS

- Processes offer big value
- Efficiency is table stakes
- Strategy affects outcomes

CHAPTER 3:

The Outcomes of Generative AI in Business Processes

Even in such early stages of implementation, 89% of survey participants are reporting measurable results from implementing AI in their business processes (Figure 3.1).

62% cited increased efficiency, the most popular response by nearly 20 points. The second ranked outcome overlaps with efficiency (cost reduction), at 43%.

Surprisingly, business leaders are also signaling that AI in business processes is leading to more powerful outcomes with long-term, compounding impacts, including:

43% have achieved **improved business agility.**

Nearly 40% are seeing **enhanced decision making.**

31% selected **faster innovation cycles.**

These long term impacts are associated with competitive advantages that compound over time, allowing them to capture unfair shares of the markets they compete in.

> "Most AI applications to date seek to improve a given task. But this is missing the larger picture; smart companies are viewing the introduction of AI as the rationale for a new look at end-to-end processes."
>
> TOM DAVENPORT, PROFESSOR AND BESTSELLING AUTHOR OF PROCESS INNOVATION: REENGINEERING WORK WITH INFORMATION TECHNOLOGY

FIGURE 3.1: WHAT MEASURABLE RESULTS HAVE YOU OBSERVED WITH GENERATIVE AI IN YOUR BUSINESS PROCESSES? (SELECT ALL THAT APPLY)

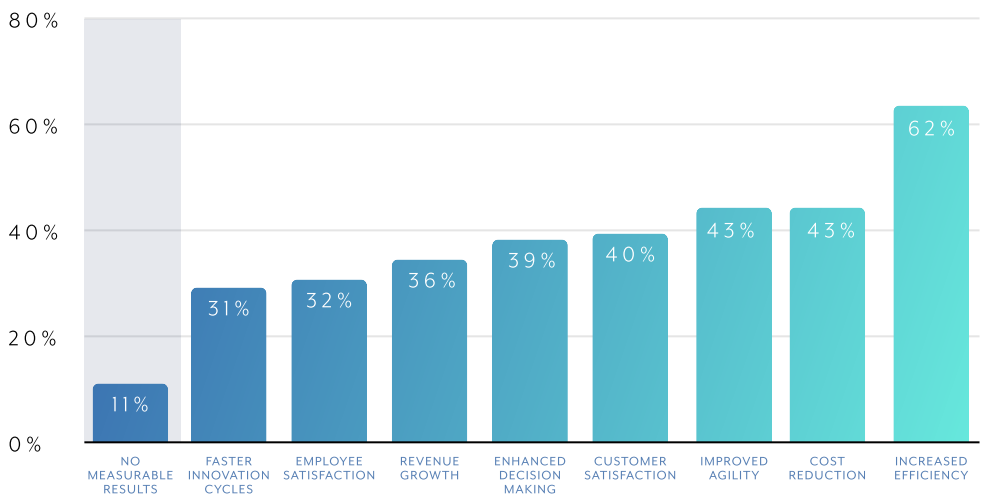

When we look at these outcomes on a departmental level - we can see that IT is consistently at the forefront of driving business outcomes - and in particular is leading in driving agility, decisioning, and innovation with AI (Figure 3.2).

FIGURE 3.2: MEASURABLE AI RESULTS BY DEPARTMENT

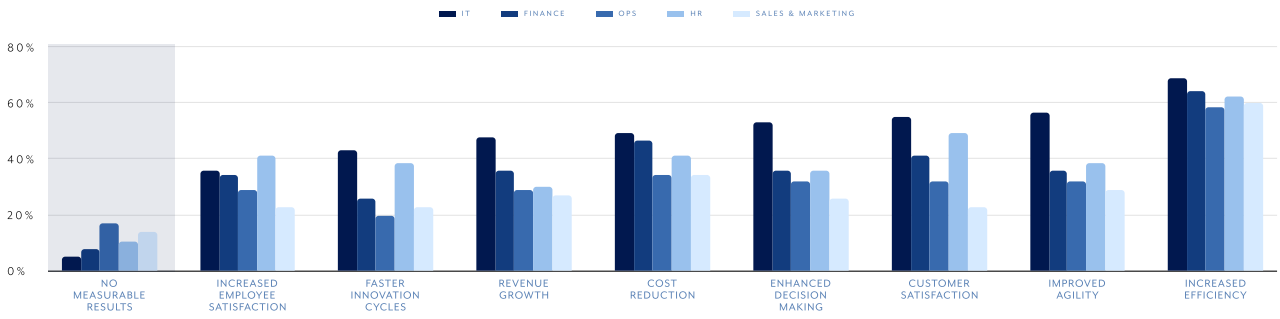

Notably, HR seems to be getting the highest increase in innovation cycles out of generative AI in their business processes. Sales and marketing seems to experience relatively less of every benefit with the exception of increased efficiency.

LLM approaches and activation strategies drive unique outcomes

Often, the choice of which LLM approach companies make is based on resources, leadership buy-in on the AI strategy, and skill. But the data shows that each approach has different strategic implications (Figure 3.3).

If, rather than efficiency, you prioritize revenue growth, enhanced decision making, and improved agility, customizing a foundational LLM with retrieval-augmented generation (RAG) is the clear choice. The improvement in enhanced decision making is quite striking, in some cases 2x higher than the other choices.

SaaS Features & Specialized LLMs are **2x more likely** to report no outcomes

RAG is nearly **2x more likely** to drive enhanced decision making, agility, and revenue growth

If average but reliable performance is the goal, a foundational LLM alone is a strong choice. It is worth noting that participants relying on foundational LLMs are the least likely to report no measurable results from their AI strategy.

The data seems to also indicate that - at this stage in generative AI technology - custom LLMs and AI features of SaaS applications are lower performing options for business processes.

FIGURE 3.3: MEASURABLE AI RESULTS BY LLM STRATEGY

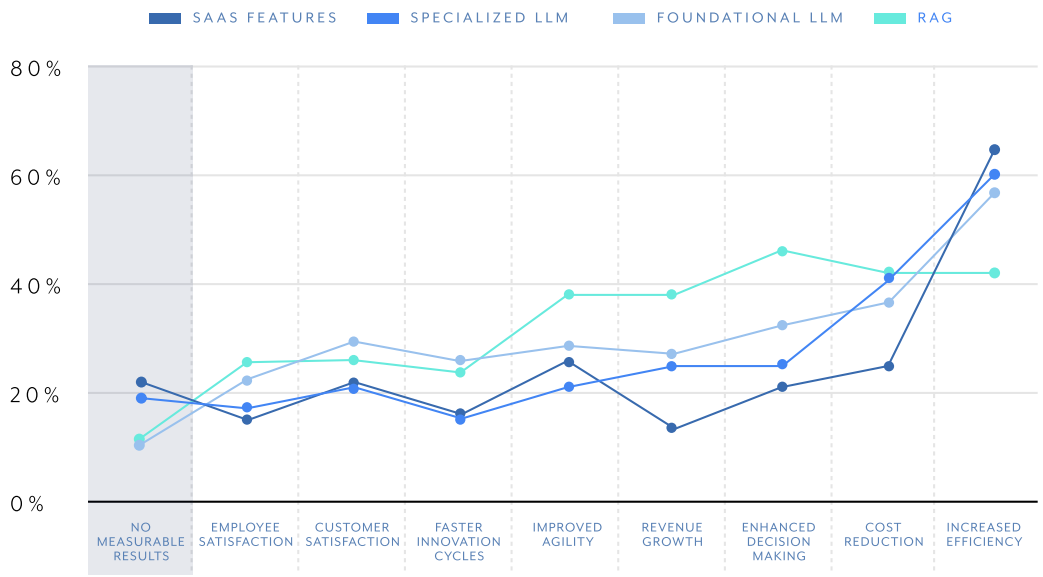

In addition, how the LLM is activated in the business process also dictates outcomes in a significant way. In chapter 2 (Figure 2.3) we saw that business leaders have a strong preference to bring LLMs into their processes with APIs. That preference may be in part due to the performance benefits - business leaders report markedly superior outcomes across the board when using APIs, and were far less likely to report no measurable results (Figure 3.4).

FIGURE 3.4: MEASURABLE RESULTS BY AI ACTIVATION APPROACH

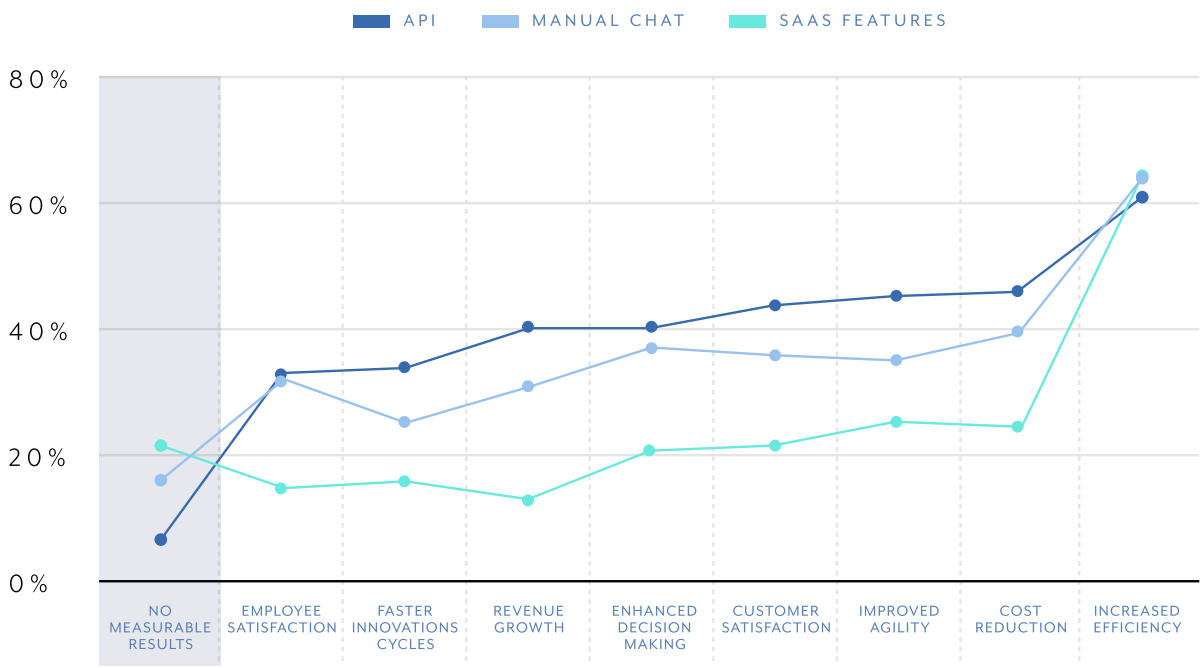

Early adopters of AI in operations have the upper hand

If generative AI is going to live up to the hype, then companies that have fully deployed AI in business processes should see better results than those who are in deployment or experimenting with the technology.

That's exactly what the data shows. Early adopters of AI in business processes are jumping into the lead.

Companies still just experimenting with AI are 4x more likely to say that they have achieved no results, whereas early adopters are nearly 2x more likely to report increased agility, enhanced decision making, faster innovation. These categories have compounding, exponential results down the road - meaning that these early adopters are outperforming their competitors today, and building competitive advantages for the future.

This explains why, in Figure 1.1, competitors are the biggest driving force behind AI adoption. Business leaders realize that this performance gap will create a new class of business winners and losers long term.

FIGURE 3.5: MEASURABLE OUTCOMES BY AI IMPLEMENTATION MATURITY

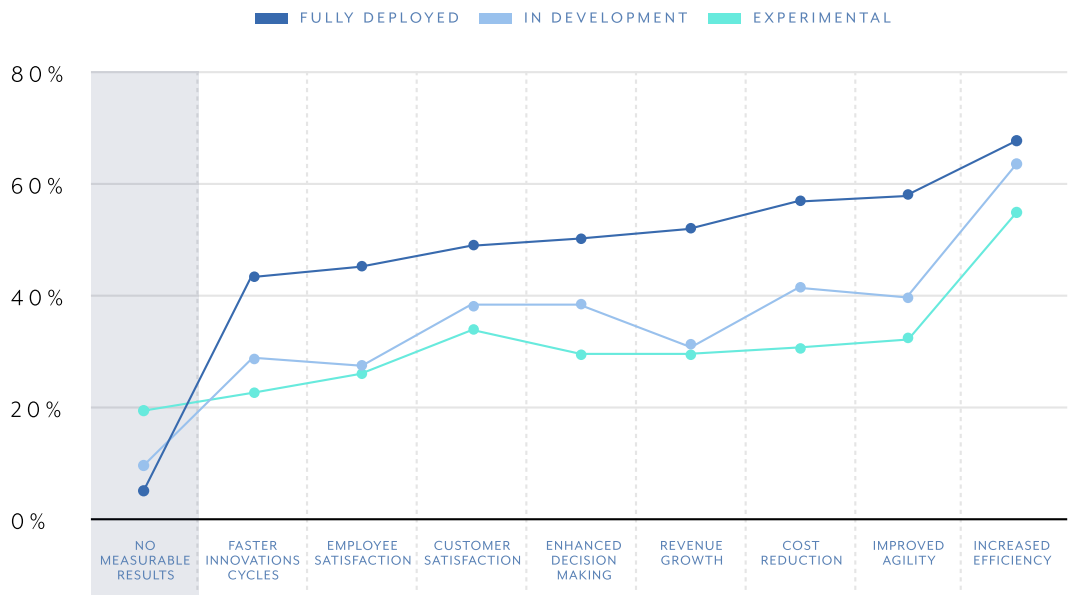

The chart shows that those experimenting with AI in their operations are falling behind those who have fully deployed a solution. However - even the experimenters are gaining serious ground on the 11% of business leaders who have not yet begun using AI in their processes.

But generative AI is not a magic pill without side effects - in fact, as several high profile failures of generative AI projects illustrate, there is a strong need for governance and human oversight. Generative AI models can hallucinate, they can demonstrate bias, and they can make mistakes. In the next two chapters, we explore these crucial topics.

KEY TAKEAWAYS

- Governance is the #1 priority
- Data security & privacy are top concerns

CHAPTER 4:

Governance of AI in Business Processes

AI presents an incredible opportunity, but without the right guardrails, it can create risk. Without governance, LLM-powered business processes can expose sensitive data, confidently make poor decisions, and hallucinate.

The survey data shows that business leaders are embracing AI with eyes wide open. Trust in LLMs is mixed - when asked how reliable and consistent outputs are from generative AI, only a third (31%) rated LLMs as "very accurate" (see Figure 5.1 in the next chapter for the full results). When asked what would improve their trust in LLMs in their business processes, the number one answer was **strong governance** (Figure 4.1)**. In fact, 90+% of participants say governance is important** for gen AI in their business processes (Figure 4.2).

FIGURE 4.1: WHAT WOULD IMPROVE YOUR TRUST IN GENERATIVE AI TO OPERATE AUTONOMOUSLY IN YOUR ORGANIZATION? (SELECT ALL THAT APPLY)

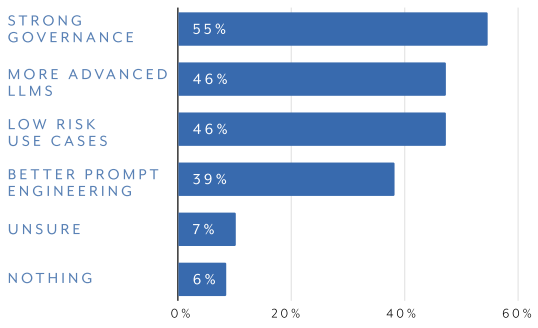

FIGURE 4.2: HOW IMPORTANT IS GOVERNANCE WITH GENERATIVE AI IN ORGANIZATIONS?

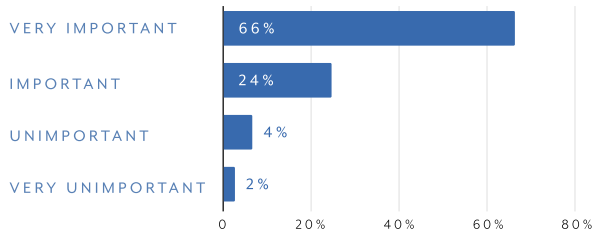

Notably - for all of the research on improving model performance, such as better data, context, and more - strong governance outranks more advanced LLMs in the minds of business leaders.

The reasons why became evident when we asked them what their top concerns are with AI in their processes. By a wide margin, governance challenges of security and privacy outrank model performance concerns such as hallucinations, loss of control and replacement (decisioning), bias, and more (Figure 4.3).

FIGURE 4.3: WHICH GENERATIVE AI TOPIC DO YOU CONSIDER CONCERNING OR WORRISOME?

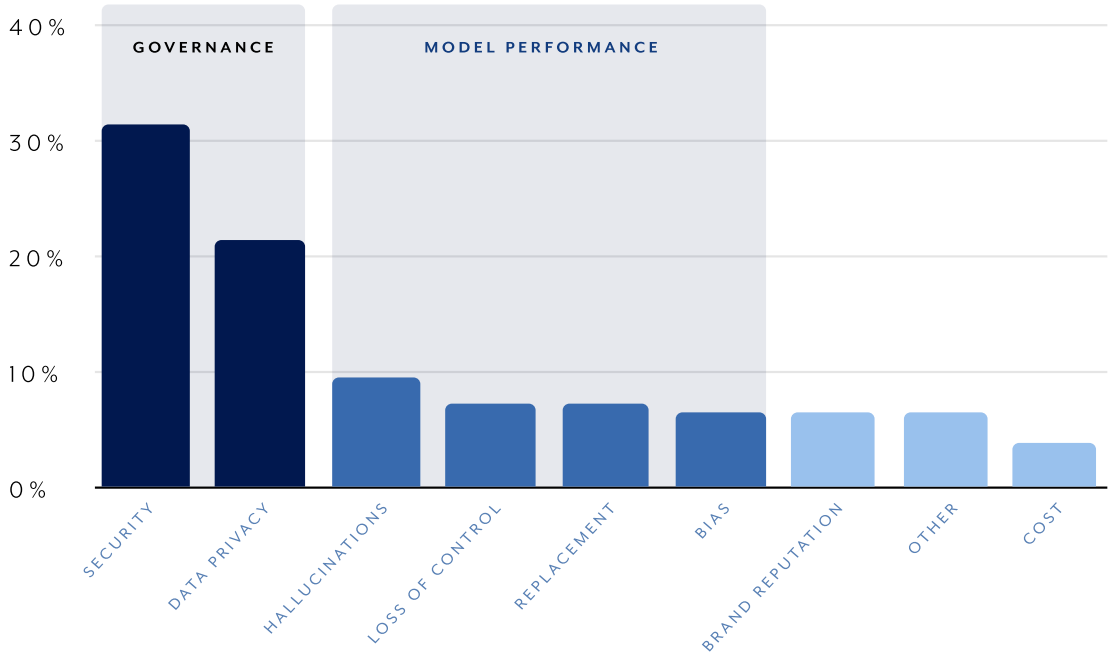

It is also important to note that, although LLM costs have been making headlines, cost ranked lowest on the list of business leader concerns. For the moment, the costs of generative AI are being subsidized by LLM vendors and their investors - this data point may change if AI vendors begin to prioritize profitability in their business models.

While all departments share similar feelings about the top three concerns, a few departments have strong opinions about select areas (Figure 4.4).

- HR, Finance, and IT Operations, who are responsible for much of the most sensitive data in the company (PII, Salary, etc), are more concerned about security and data privacy.
- Sales and marketing, who own the messaging to customers and prospects, are more concerned about hallucinations.

FIGURE 4.4: TOP 3 AI CONCERNS BY DEPARTMENT

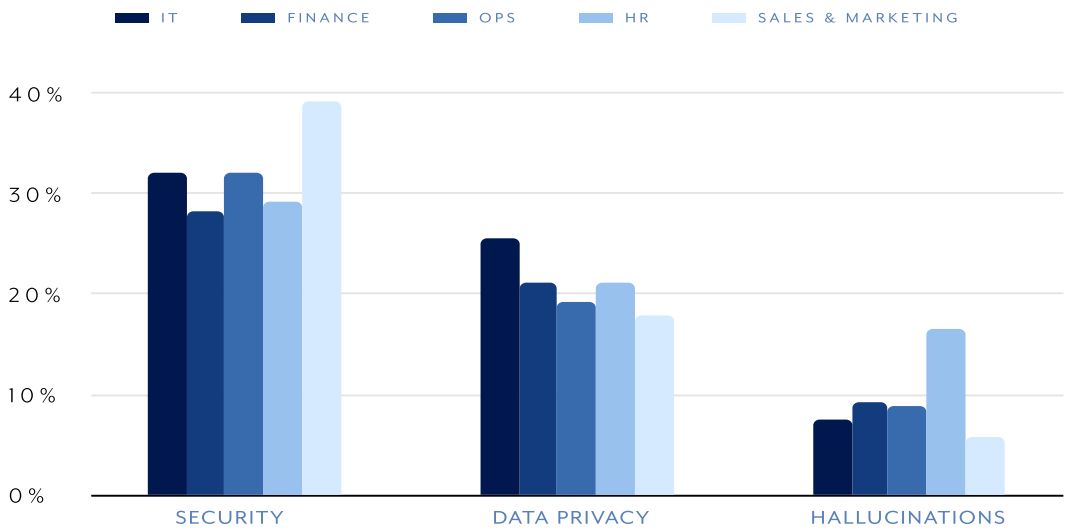

AI governance isn't just a protective measure for organizations. It's a proactive strategy to enable employees to use the technology more effectively, from upholding the privacy of proprietary data to establishing ethical standards to prevent reputational fallout. When we asked participants which areas of governance they are prioritizing, security practices such as data privacy and protection comes out on top (75%), followed by ethical guidelines and practices (59%), and access controls (53%) (Figure 4.5).

FIGURE 4.5: WHICH TYPES OF GENERATIVE AI GOVERNANCE ARE YOU PRIORITIZING?

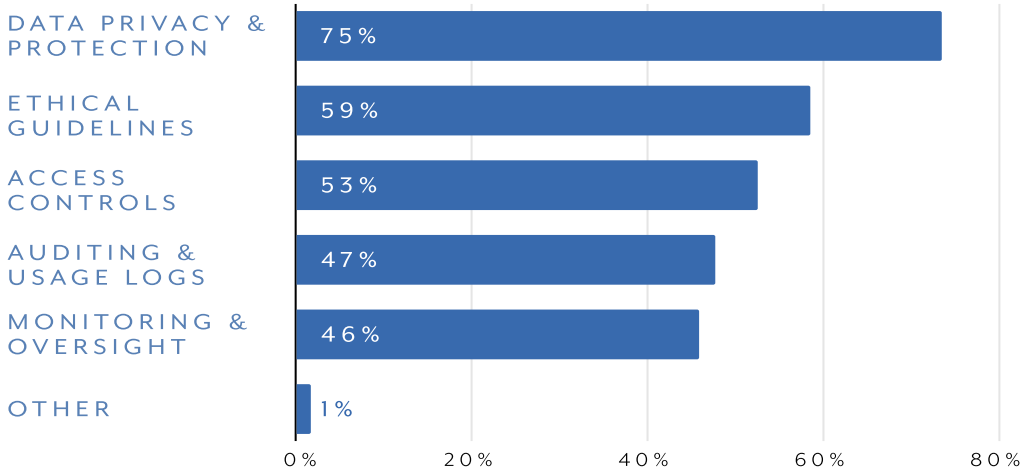

With AI governance in place to guide ethical, safe, and responsible use, human oversight is essential to keep AI decisions aligned with company goals. In the next chapter, we'll look at how people help monitor and steer AI to ensure it works as intended.

KEY TAKEAWAYS

- People have a big role in an AI world
- Human expertise will grow in value
- Leaders balance autonomy and oversight

CHAPTER 5:

The Role of People in a World of Autonomous AI and Agents

As mentioned in the previous chapter, business leaders have misgivings about generative AI hallucinating the wrong answer. When asked about the reliability and consistency of Gen AI model outputs, 47% say outputs are "somewhat accurate," and only 31% say very accurate (Figure 5.1).

FIGURE 5.1: IN YOUR OPINION, HOW RELIABLE AND CONSISTENT ARE THE OUTPUTS FROM THE CURRENT GEN AI MODELS?

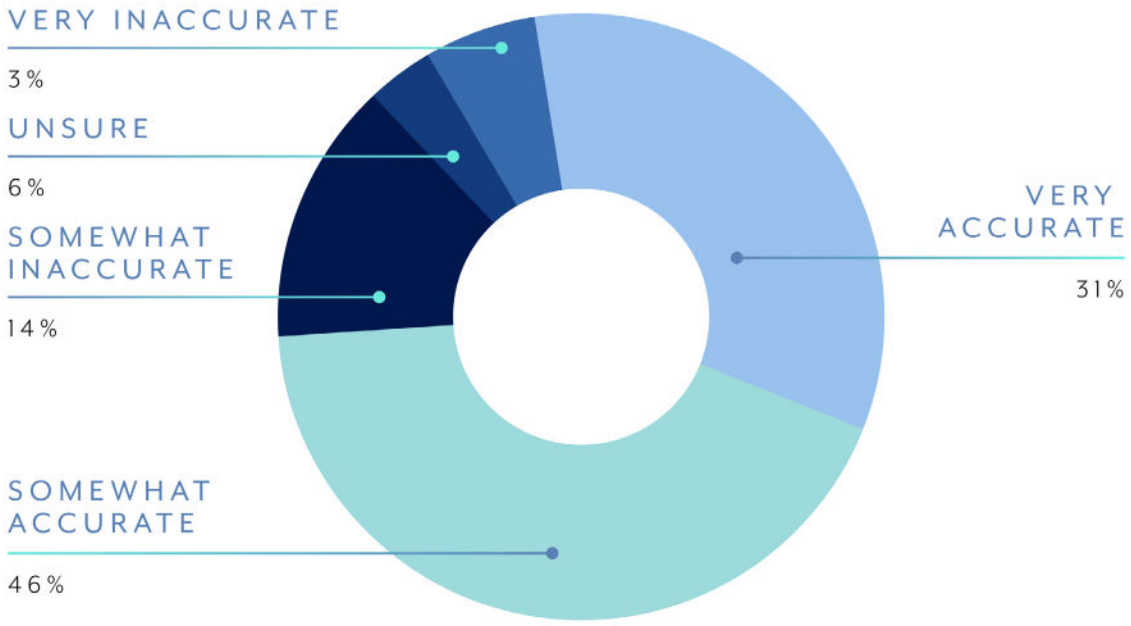

In a business process, how much inaccuracy is acceptable? How often should a person step in to review, revise, and approve AI decisions and content?

In February 2024, Air Canada refused to honor a refund offer hallucinated by the airline's AI chatbot. If an employee had been in the loop to review any offers the chatbot made, the airline might have avoided the resulting lawsuit and reputation damage after the lawsuit made international headlines.

When asked how they felt about human review - participants shared an overwhelming mandate - 94% said that human oversight and approvals of generative AI outputs are important (Figure 4.2), and 95% believe that human oversight and approvals are important when working with autonomous AI agents (Figure 5.2).

FIGURE 5.2: IN YOUR OPINION, HOW IMPORTANT IS HUMAN OVERSIGHT AND APPROVALS OF GENERATIVE AI OUTPUTS?

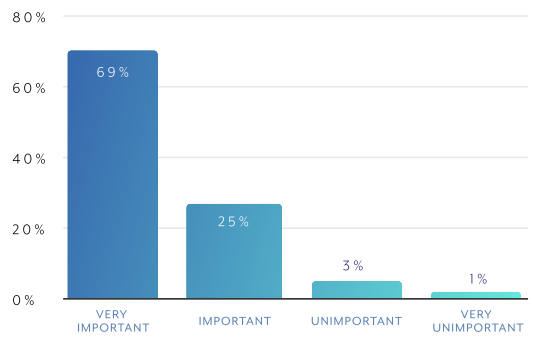

FIGURE 5.3: WHEN WORKING WITH AI AGENTS, IN YOUR OPINION HOW IMPORTANT ARE HUMAN OVERSIGHT AND APPROVALS (HUMAN IN THE LOOP?)

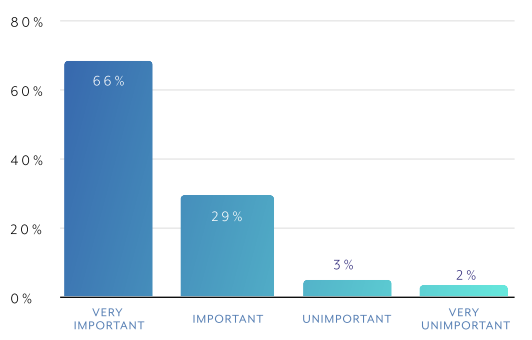

"AI has some challenges. It can be wrong, it can hallucinate, it can use copyrighted materials, it can be biased.

The antidote is to put a human in the loop. If we are thinking of collaborative AI, and thriving with AI, the human in the loop becomes a necessary part of that equation. Then, the human takes on responsibility for the tool.

I often talk about the wielder of a hammer. Did the hammer break the glass, or did the person who wielded the hammer? It becomes the responsibility of the wielder of the tool to ensure that that tool is used in the most effective way possible. That's why a human must be in the loop to take the maximum advantage and keep AI responsible.

If you can do that, keep the human in the loop, then I think we can unleash the full power of AI.

EARL NEWSOME, CIO, CUMMINS INC.

Is it a total rejection of autonomous AI if 95% of business leaders believe humans in the loop are important? No. Nearly 60% of business leaders are comfortable with autonomous AI decisions and content sent to customers (Figures 5.4 and 5.5).

FIGURE 5.4: HOW COMFORTABLE ARE YOU WITH GENERATIVE AI SENDING CONTENT AND MESSAGES AUTONOMOUSLY TO CUSTOMERS WITH NO HUMAN REVISIONS AND APPROVALS?

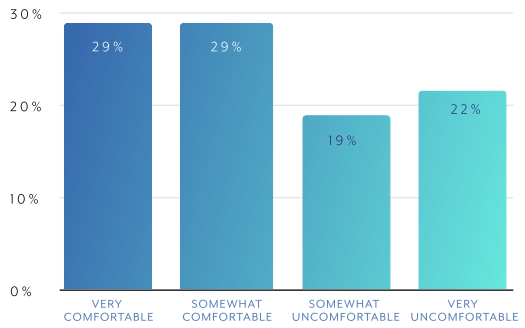

FIGURE 5.5: HOW COMFORTABLE ARE YOU FOR GENERATIVE AI TO MAKE DECISIONS AUTONOMOUSLY, WITH NO HUMAN OVERSIGHT AND APPROVALS IN YOUR ORGANIZATION?

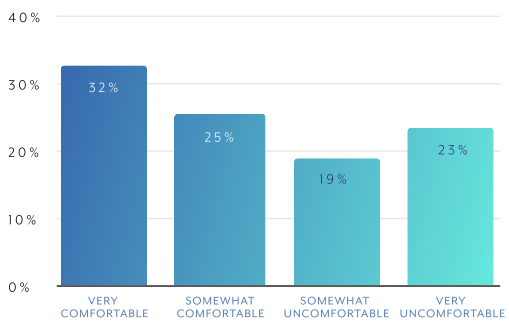

In practice, how comfortable business leaders will be with autonomous AI will likely depend on a case by case basis. The level of trivial vs. sensitive content, and low level operational vs. strategic decisions will influence how often AI is let loose with minimal human oversight.

"We are firm believers, particularly with generative AI, that the human is in the loop. You need to make sure that you make the decisions...You may have a situation where generative AI makes up answers, or hallucinates.

If you think about generative AI to generate some artwork. Well, art is in the eye of the beholder. That's probably okay.

But, in the field of hiring, for example, that's not okay. If it is generating answers or selecting a match between a person and a job, and it gets it wrong, if it doesn't have a way of explaining itself, that technology is not appropriate."

AL SMITH, CHIEF TECHNOLOGY OFFICER, ICIMS

90%+ of respondents say at least half of business processes should not be fully autonomous - and should include humans in the loop (Figure 5.6).

FIGURE 5.6: HOW MUCH OF GENERATIVE AI DECISIONS REQUIRE HUMAN OVERSIGHT BEFORE BEING USED IN A BUSINESS CONTEXT?

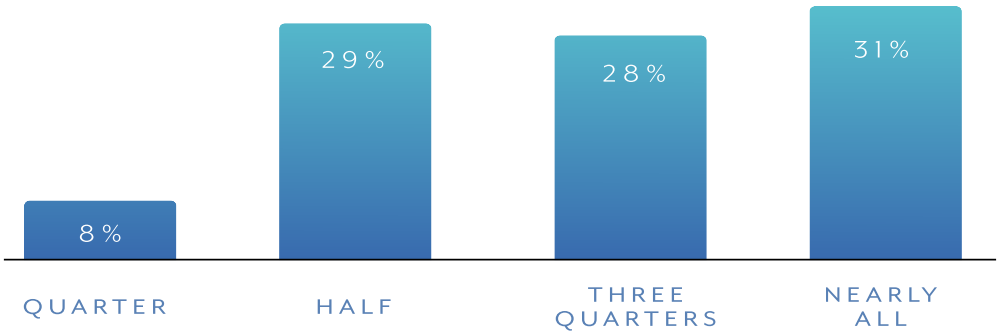

Evidence suggests that the preference for humans in the loop is about more than responsible AI. It is also correlated with better performance. Recent research by MIT and Accenture highlighted the benefits of nudging humans to review generative AI outputs. This form of human-AI augmentation led to measurably better performance:

"The findings revealed that consciously adding some friction to the process of reviewing LLM-generated content can lead to increased accuracy — without significantly increasing the time required to complete the task. This has implications for how companies can deploy generative AI applications more responsibly."

90% of business leaders believe AI should have a human in the loop for at least half of all processes

Who should review?

Lastly - ***who*** should be responsible for conducting this review and approval? Conversations about the "autonomous enterprise" have sparked fears that companies will operate like fast food restaurants: Automation and AI handle most of the work, and the company is staffed by entry-level employees.

However, a recent article from Wharton researchers suggested a counterintuitive hypothesis: generative AI is going to increase demand for expertise, rather than decrease it: "One challenge when integrating LLM output with human oversight is that in many contexts, the human must know something about the domain to be able to assess whether the LLM output is valuable. This suggests that specific knowledge cannot be "outsourced" to an LLM — domain experts are still needed to evaluate whether LLM output is any good before it is put into use."

The survey data reinforces this counterintuitive point - that generative AI in business processes will increase demand for human expertise. Overwhelmingly, business leaders believe that managers and subject matter experts should be tapped to review generative AI outputs and decisions, rather than front line or entry level employees (Figure 5.7).

FIGURE 5.7: FOR MOST CASES, WHAT LEVEL OF HUMAN OVERSIGHT AND REVIEW WOULD YOU CONSIDER APPROPRIATE FOR GENAI?

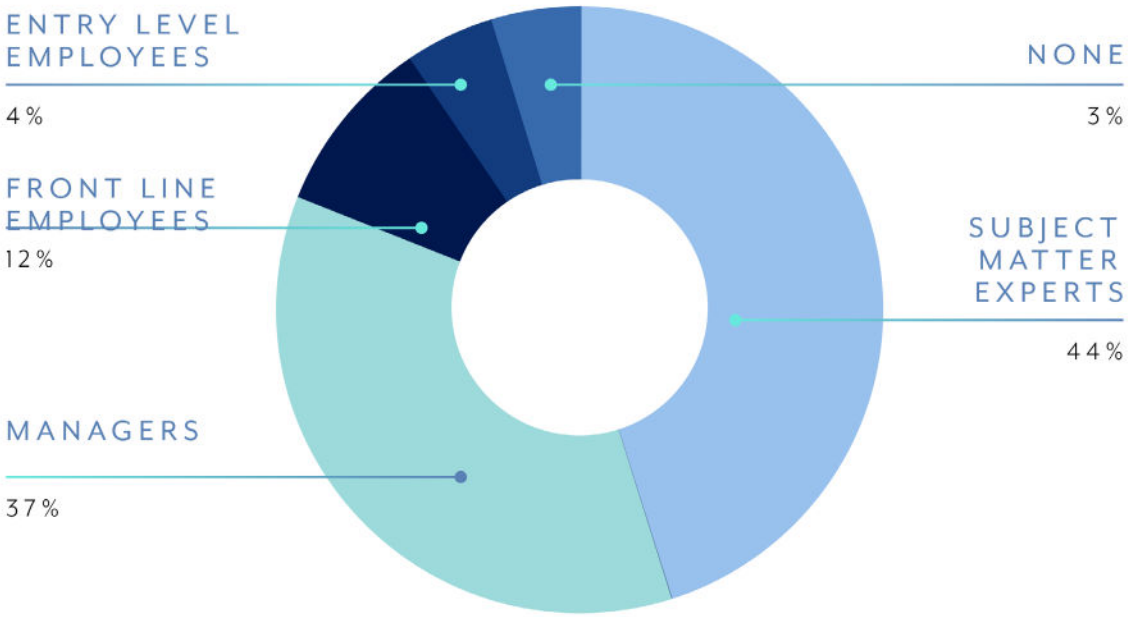

Many are still getting used to the idea of reviewing human-like text generated by a machine. But the next wave of transformation is around the corner in the concept of AI agents. In the next chapter, we'll take a closer look at this exciting topic.

CHAPTER 6:

AI Agents & The Future of Business Processes

In an article titled *Why Agents are the Next Frontier of Generative AI*, McKinsey predicts "We are beginning an evolution from knowledge-based, gen-AI-powered tools—say, chatbots that answer questions and generate content—to gen AI–enabled "agents" that use foundation models to execute complex, multistep workflows across a digital world. In short, the technology is moving from thought to action."

Whether business leaders are ready or not, agentic processes are around the corner. Agentic processes are workflows that are partially or completely executed by an AI agent. Most, but not all business leaders are ready for AI agents to handle work on their behalf (Figure 6.1).

FIGURE 6.1: IN THE FUTURE OF WORK, HOW DO YOU FEEL ABOUT THE IDEAS OF AI AGENTS COMPLETING WORK ON YOUR BEHALF?

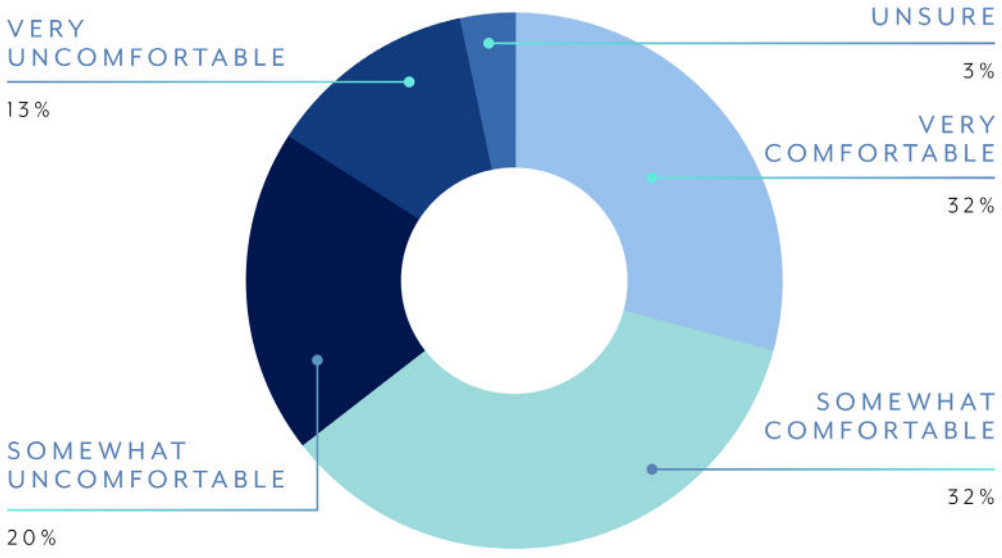

Agentic readiness varies by department. IT and Finance leaders report feeling more comfortable with the idea, whereas Sales & Marketing leaders are less comfortable (Figure 6.2). It is possible that departments like Finance and IT have been subjected to more automation over the years, leading them to be more optimistic about AI agents than others.

FIGURE 6.2: IN THE FUTURE OF WORK, HOW DO YOU FEEL ABOUT THE IDEAS OF AI AGENTS COMPLETING WORK ON YOUR BEHALF? (COHORT FROM FIGURE 0.1)

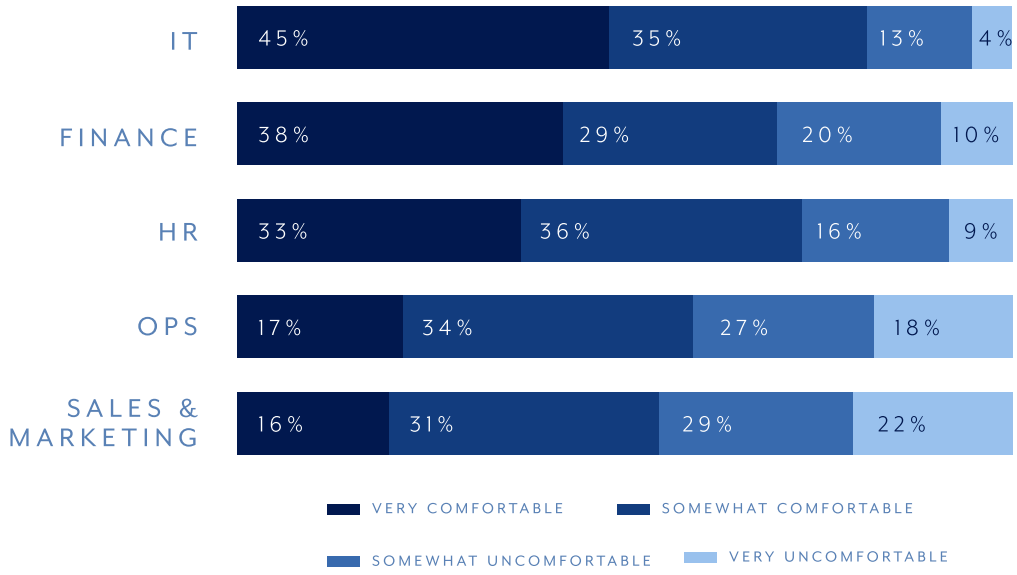

Many predict there will be armies of agents in every company, assisting every employee in their roles. They will empower employees with superhuman work capacity, leading to what SiliconAngle Research describes as "5-10x improvements in cognitive labor productivity."

"When we think about automation and AI, we think of it more as an augmentation tool versus something that replaces humans."

DARRYL MARAJ, CIO, TBC CORPORATION

The business leaders in this survey seem to agree. When asked which agentic outcome they were most excited about - productivity and efficiency rose to the top (Figure 6.3).

FIGURE 6.3: WHEN CONSIDERING AI AGENTS, WHICH SOUNDS MOST EXCITING TO YOU?

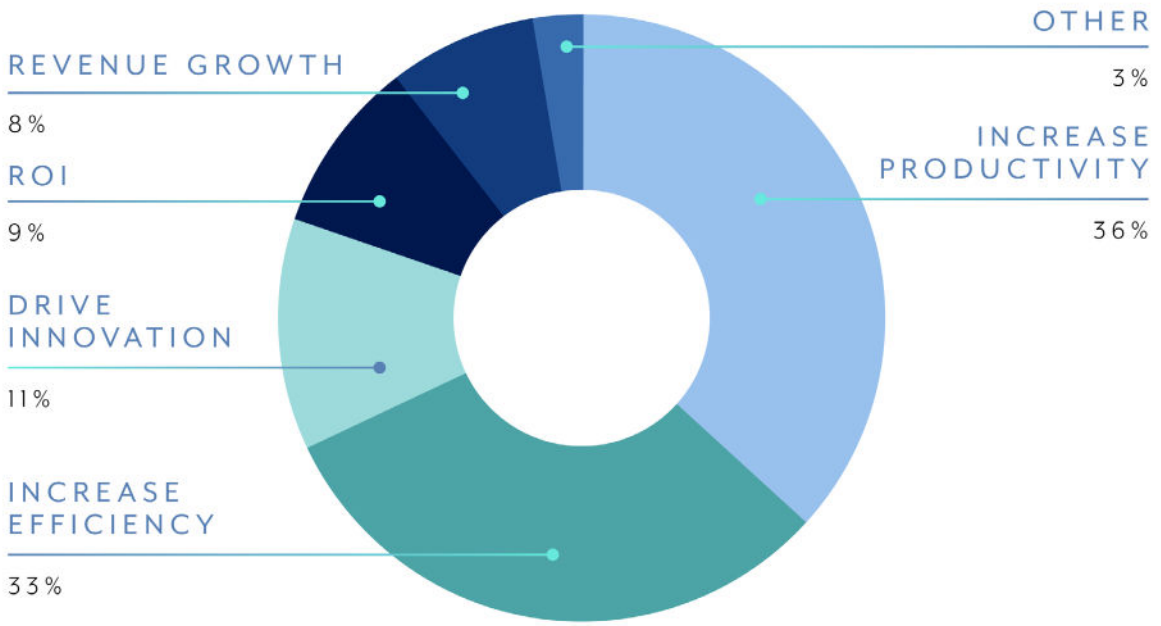

What should today's outcomes teach us about tomorrow's agents?

An AI agent can:

- Generate plans, make decisions, and take action across systems
- Consult employees for review and approval
- Collaborate with other agents to execute complex processes
- Learn over time to improve
- Operate in a responsible, governed way

An LLM deployed by itself is not an agent - it is just a smart chat bot, and a liability. Consider the Air Canada example that we mentioned earlier. The chat bot did not have the context it needed - so it hallucinated a refund that violated Air Canada's own policies. And it was not empowered to take the action that it needed, so it was not able to run its decision by a human agent for approval.

In these early days of agentic technology, most leaders are too early in their agentic journey to report meaningful results. However, there are lessons to be learned from the results leaders are seeing by generally adding AI to their business processes. After all, these are the predecessors to the agentic future. What can they teach us?

The power of context

We can see the power of context when we look at how retrieval-augmented generation (RAG) impacts results (the charts below were originally featured in chapter 3). The RAG method allows the LLM to leverage data that would not be in its original training data, such as the latest customer data, employee data, or other. In the Air Canada example, an agent might want to access the latest refund policies, the passenger's account data, and other available flights.

The value of having the latest data - which we can see in the marked performance improvement on decisioning, revenue growth, and agility - will only grow in the agentic future:

THE POWER OF CONTEXT FOR AI AGENTS (FIGURE 3.3)

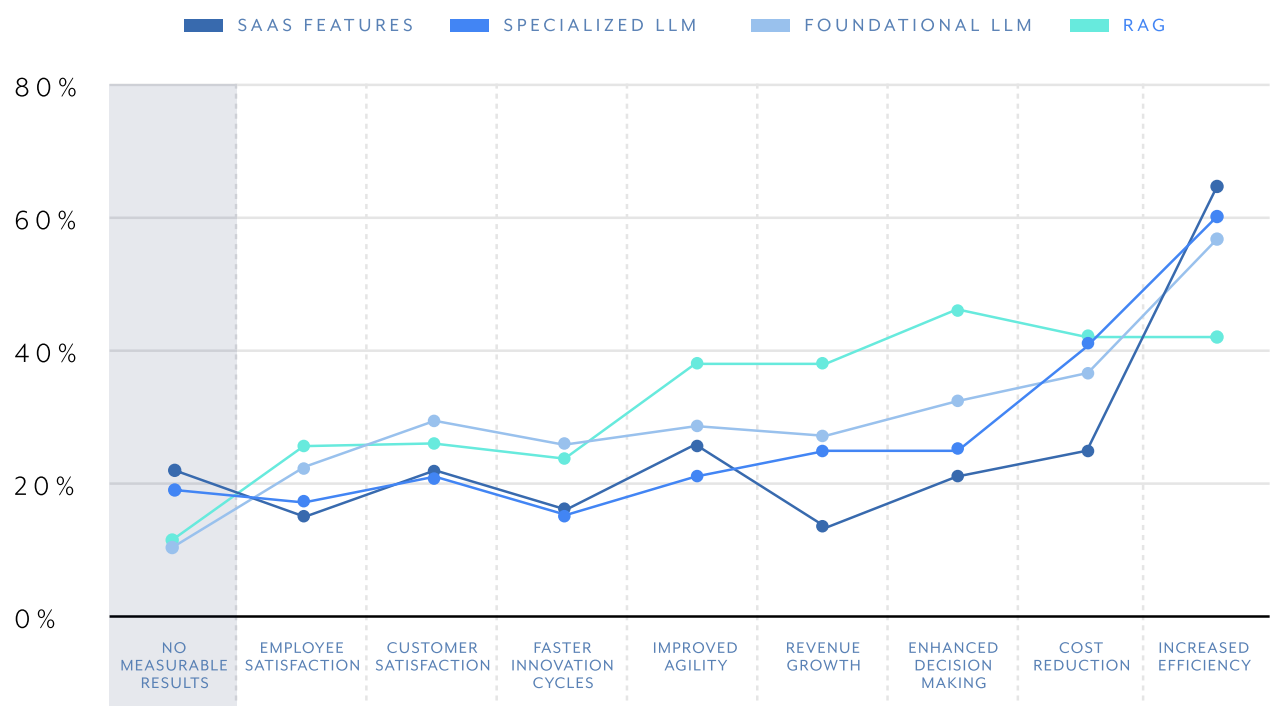

The ability to take action

As the McKinsey quote earlier in this chapter notes, agents are about AI technology "moving from thought to action." The ability to take action across systems is paramount for agents to succeed. A sales agent would want to be able to take action across customer relationship management (CRM), a sales engagement, and sales enablement platforms.

The strong preference we see in today's data for API-based usage of LLMs (Figure 2.3) and enhanced performance of LLMs in business processes for those leveraging API-based approaches (Figure 3.4) strongly suggests that agents empowered to take action across API-based connectors will outperform others.

The agentic gap

Early adopters of agentic technologies will enjoy similar benefits to the early adopters of LLMs in business processes. As we saw in chapter 3, the more mature the AI implementation, the greater the reported business value. However, the difference is that agentic technology will, as SiliconAngle reports, confer 5-10x gains in productivity upon employees. Multiplied by hundreds or thousands of employees across the company, this will undoubtedly create a major "agentic gap" between agentic and traditional companies. The agentic gap will resemble, but be far wider than the gap we see today:

FIGURE 3.5: MEASURABLE OUTCOMES BY AI IMPLEMENTATION MATURITY

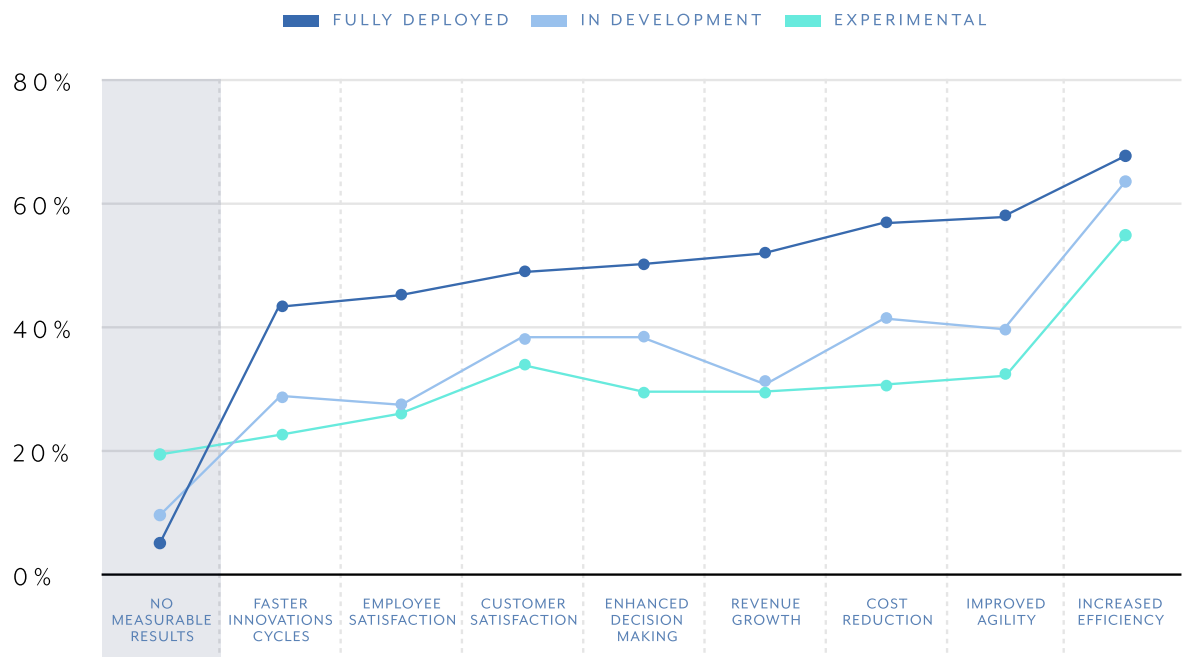

The lesson for leaders: Ignore the agentic revolution at your peril. Despite the headlines, despite the skepticism - the evidence shows that the early adopters of generative AI in business processes are building an advantage over those who are not.

Conclusion

Recent reporting from the Census Bureau shows that only 4.8% of American companies use AI to produce goods and services. Many who question the hype around generative AI, including The Economist, cite this statistic. But questioning where the AI-generated products are, or why AI features aren't driving up SaaS revenues, or other theories may be missing the point.

As this survey revealed - the real value, at least in these early days of generative AI innovation - is likely being felt in operations. In their business processes. Here, companies expect to gain efficiency and save costs. But the data also shows that companies are building competitive advantages - such as business agility, better decision making, and more.

To maximize the value of generative AI in business processes, business leaders are grappling with big strategic questions.

- How do we activate AI in our operations?
- Which approach(es) do we take with LLMs?
- How often and when do we include humans in our AI-powered processes?
- What do we need to govern to make sure this innovation does not create risk?

The evidence is clear - despite AI doom in the press - that business leaders are driven by much more than market sentiment to explore these questions. The threat of their competitors leaping ahead in the AI arms race is enough to drive them forward. As we enter a new agentic economy, the threat of losing out on an agentic gap will only motivate them further.

It is up to business leaders to ensure that their organizations make the most of this disruptive moment - or risk being left behind.

Methodology & Demographics

Workato commissioned an independent market survey from UserEvidence of over 1,000 business leaders in North America about generative AI in their business processes. The research sample was vendor-neutral, and did not target Workato or UserEvidence customers, although they were not excluded from participating.

The largest cohort (31%) worked within their company's technology or IT department. Respondents also worked in departments like sales & marketing (15%), finance (10%), operations (10%), and human resources (8%) as illustrated in figure 0.1.

FIGURE 0.1: WHAT DEPARTMENT DO YOU WORK IN?

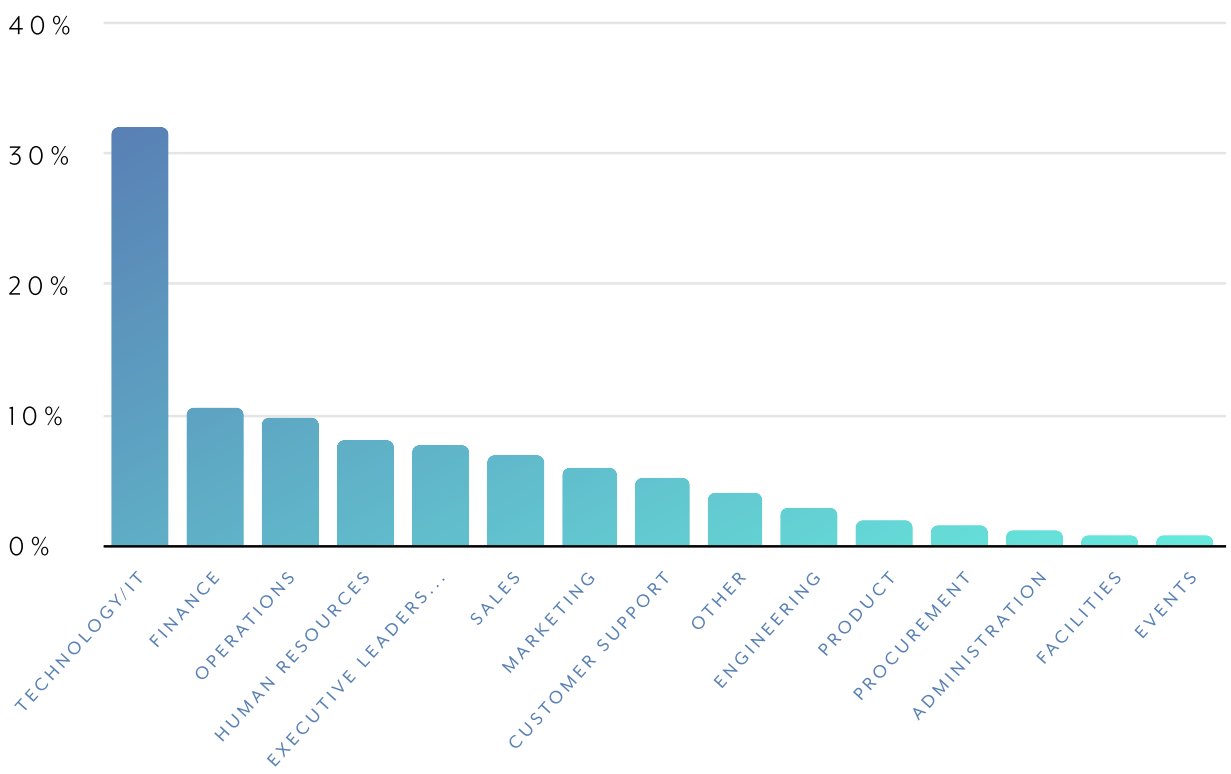

More than half (57%) worked at the director level, 26% said they worked at the C-suite level, and the remaining 16% said they were at the VP level. (Figure 0.2)

FIGURE 0.2: WHAT BEST DESCRIBES YOUR ROLE WITHIN THE COMPANY?

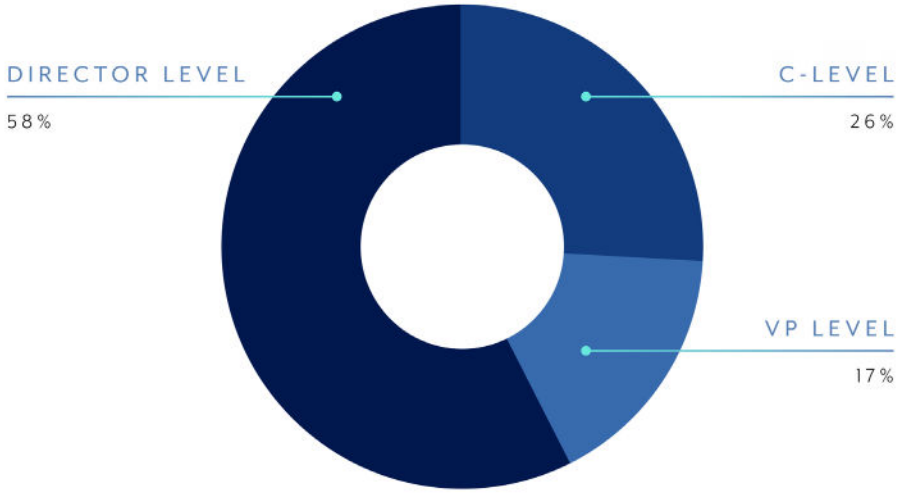

Respondents represented nearly 20 industries, but nearly half came from four key sectors: Technology, financial services, healthcare, and manufacturing (see figure 0.3):

FIGURE 0.3: WHAT INDUSTRY DO YOU CURRENTLY WORK IN?

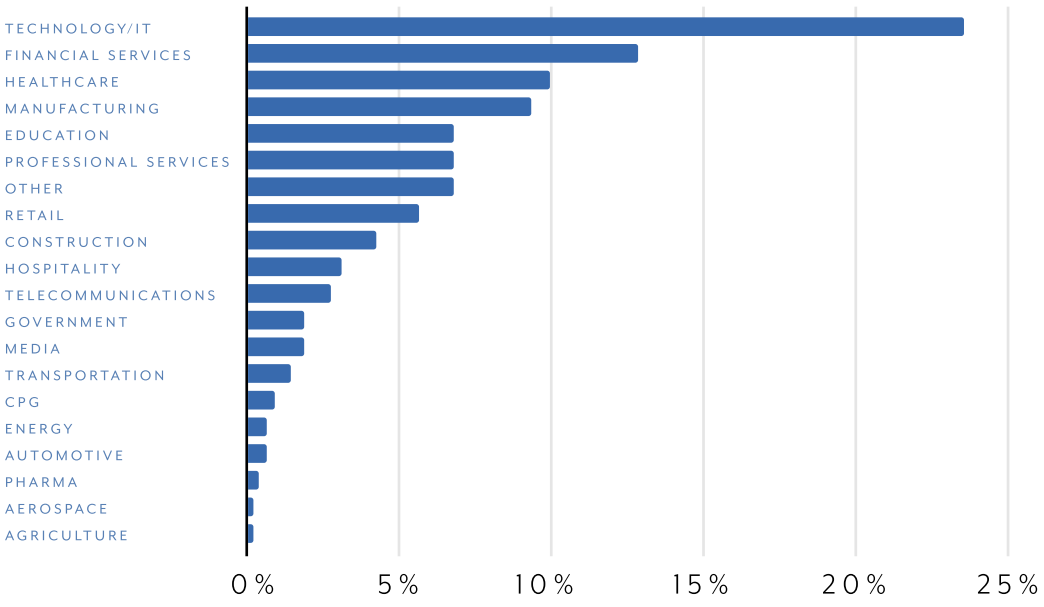

The majority of respondents came from mid-market or enterprise companies (see Figure 0.4):

- **15%** work at a company with **250-499** employees
- **44%** work at a company with **500-2,499** employees
- **40%** work at companies with **2,500** employees or more
- **20%** work for companies with **5,000** companies or more
- Almost **15%** work for companies with **10,000+** employees

FIGURE 0.4: HOW MANY EMPLOYEES WORK AT YOUR COMPANY?

About UserEvidence

UserEvidence is a software company and independent research partner that helps B2B technology companies produce original research content from practitioners in their industry. All research completed by UserEvidence is verified and authentic according to their research principles: Identity verification, significance and representation, quality and independence, and transparency. All UserEvidence research is based on real user feedback without interference, bias, or spin from our clients.

UserEvidence Research Principles

These principles guide all research efforts at UserEvidence—whether working with a vendor's users for our Customer Evidence offering, or industry practitioners in a specific field for our Research Content offering. The goal of these principles is to give buyers trust and confidence that you are viewing authentic and verified research based on real user feedback, without interference, bias, and spin from the vendor.

Principle 1 — Identity verification.

In every study we conduct, UserEvidence independently verifies that a participant in our research study is a real user of a vendor (in the case of Customer Evidence) or an industry practitioner (in the case of Research Content). We use a variety of human and algorithmic verification mechanisms, including corporate email domain verification (i.e., so a vendor can't just create 17 Gmail addresses that all give positive reviews).

Principle 2 — Significance and Representation.

UserEvidence believes trust is built by showing an honest and complete representation of the success (or lack thereof) of users. We pursue statistical significance in our research and substantiate our findings with a large and representative set of user responses to create more confidence in our analysis. We aim to canvas a diverse swatch of users across industries, seniorities, personas — to provide the whole picture of usage and allow buyers to find relevant data from other users in their segment, not just a handful of vendor-curated happy customers.

Principle 3 — Quality and Independence.

UserEvidence is committed to producing quality and independent research at all times. This starts at the beginning of the research process with survey and questionnaire design to drive accurate and substantive responses. We aim to reduce bias in our study design and use large sample sizes of respondents where possible. While UserEvidence is compensated by the vendor for conducting the research, trust is our business and our priority, and we do not allow vendors to change, influence, or misrepresent the results (even if they are unfavorable) at any time.

Principle 4 — Transparency.

We believe research should not be done in a black box. For transparency, all UserEvidence research includes the statistical N (number of respondents), and buyers can explore the underlying blinded (de-identified) raw data and responses associated with any statistic, chart, or study. UserEvidence provides clear citation guidelines for clients when leveraging research, including guidelines on sharing research methodology and sample size.