



IAVOR I. BOJINOV
KARIM R. LAKHANI
DAVID LANE

JPMorganChase: Leadership in the Age of GenAI

While we do not know the full effect or the precise rate at which AI will change our business – or how it will affect society at large – we are completely convinced the consequences will be extraordinary and possibly as transformational as some of the major technological inventions of the past several hundred years: Think the printing press, the steam engine, electricity, computing and the Internet, among others.

— Jamie Dimon, JPMorganChase Chairman & CEO¹

In March 2025, Lori Beer, Mary Erdoes, and Teresa Heitsenrether were discussing the impact and potential of generative artificial intelligence (GenAI) for the roughly 4,000 advisors who served the high-net-worth clients of JPMorganChase's (JPMC) Global Private Bank (PB). Private bankers offered advice and recommendations to help clients preserve, protect, and grow their wealth, including estate planning and investment management, tax advice, and sophisticated personal banking services.

All three senior executives sat on JPMC's 15-member operating committee (see **Exhibit 1**) – Beer as global chief information officer (CIO), Heitsenrether as global chief data and analytics officer (CDAO), and Erdoes as CEO of JPMC's Asset & Wealth Management business, which included the PB. All three, together with Chief Risk Officer Ashley Bacon, had been deeply involved in the adoption, use, and governance of GenAI at JPMC since OpenAI's world-changing November 2022 release of ChatGPT, a large language model (LLM).

After first blocking ChatGPT to protect its data security, JPMC had moved steadily to embrace LLMs in ways that preserved operational reliability, accuracy, security, and compliance. By early 2024, JPMC had facilitated GenAI adoption by enabling application programming interface (API)^a access to LLMs through large public cloud providers. This approach applied JPMC security and data protection protocols for safe and reliable LLM interactions with the bank's data, most of which was stored in the public and private cloud. The bank's GenAI toolkit, LLM Suite, was on the desktops of more than 200,000 employees, half of whom used it several times a day. Connect Coach (Coach) and a plethora of other GenAI applications were generating enthusiasm by simplifying and accelerating completion of otherwise tedious and time-consuming "no joy" tasks (see **Exhibit 2**). Ongoing initiatives were applying GenAI in an effort to optimize processes at the bank, and existing AI initiatives for fraud

^a APIs connected computers and software programs, allowing them to interact.

Professors Iavor I. Bojinov and Karim R. Lakhani and Lead Case Researcher David Lane (Case Research & Writing Group) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2025 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

prevention, service personalization, trading, operational efficiencies, and credit decisions already brought JPMC up to \$1.5 billion in value.²

Coach tapped JPMC content and certain client data, selectively augmented with public information, to give PB advisors information to support the needs of their high-net-worth clients. While advisors prided themselves on high-touch, long-term client relationships and leveraged Coach to provide more timely information to their clients, Erdoes was confident that Coach also could help advisors perform better, more consistently, and much more productively than ever before.

AWM was now working to hyper-personalize Coach to available client data – from risk appetite to investment approach. Personalization via GenAI would make advising scalable, helping AWM advisors operate at the level of its most effective advisors. Erdoes also knew that some clients preferred self-service, whether for online trading or research. With the growing momentum of GenAI across the business and more broadly in society, the aim was to maximize the impact for clients and enhance JPMC's competitive advantage. Perhaps the time had come to press for a step that she had long mulled. To Beer and Heitsenrether, she now asked, "Should we consider releasing a version of Coach to our clients? If so, when? What should it look like? And how would we know when we are ready? Or should we take an intermediate step, and augment advisors' client meetings with voice-capable GenAI that can listen and answer questions?"

JPMC Business Segments

JPMC was, on multiple measures, the world's leading bank, and by far the largest by market capitalization.³ In October 2024, an AI market research firm ranked JPMC first among 50 leading banks for AI capabilities on metrics including talent and talent development, innovation, and transparency.⁴ JPMC's AI workforce was the largest among them too, larger than the next seven banks' combined.⁵

In 2024, JPMC operated three business segments: Asset & Wealth Management (AWM), Commercial & Investment Bank (CIB), and Consumer & Community Banking (CCB). Other functions were subsumed under a Corporate category (see **Exhibit 3** for selected data and metrics).

AWM comprised two sub-businesses. Asset Management offered investment products and advice to institutional and retail clients. The PB offered advice, planning, and products to clients with net worth of typically \$10 million or more. These included brokerage, custody, deposits, estate planning, lending, investment management, and retirement products and services. At the start of 2025, AWM client assets amounted to \$5.9 trillion; AWM assets under management were \$4 trillion.⁶

CIB served institutional clients in all major capital markets with banking and payments services. It also offered debt and equity raising, loan origination and syndication, and advice on capital strategy and structure. CIB's market and securities services involved trading and market-making across all main product classes. This included prime brokerage, risk management, clearing, and global custody services, as well as liquidity, data, and fund services.⁷

CCB represented JPMC's consumer and retail banking business, operating under the Chase brand via branches, ATMs, and online. CCB also issued credit cards and loans for autos and home purchases.⁸

Adopting GenAI

Early Decisions

Reactions within JPMC to the release of ChatGPT ranged from exhilaration to concern. Chief Analytics Officer Derek Waldron, for instance, saw GenAI as transforming the state of the art. “All of a sudden you have an out-of-the-box technology that applies AI to a whole new set of domains beyond traditional AI,” he said. “In the past at J. P. Morgan, to get an AI idea implemented, you would first need to prioritize it, get it funded, get an engineering team, and get it implemented. Since that took a year and a lot of effort, many ideas didn’t make it to implementation. GenAI could allow businesspeople to solve their problems without any of that.”

Bacon found GenAI “instantly amazing.” He sent word of ChatGPT to the JPMC Operating Committee on the day he tried it, urging members to set up accounts and play with it on their personal devices. “Efficiency and efficacy plays were both immediately obvious,” he said. “But my first concern was caution that accounts created within our firewall would be dumping confidential information into it. The first thing we did was to block it, while working in parallel to get an internally developed version into the hands of every employee to unleash its power.”

JPMC was a market leader in a highly regulated and highly competitive industry, maximizing for financial performance while complying with domestic and international legal and regulatory requirements. Ensuring operational reliability was imperative, too: each day, JPMC handled \$10 trillion in transactions.

Concern over GenAI as an unproven technology therefore prompted careful consideration of how best to harness its possibilities. While data security came first, it was not the only issue. Heitsenrether said, “Since data is our key differentiator, we had to ensure that our data was not being used to train external models. We also had to design a system to comply with still-evolving regulatory requirements across jurisdictions.”

JPMC’s solution was to build an internal platform and receive LLM access via Microsoft Azure’s cloud computing services. Microsoft was OpenAI’s exclusive provider of computing power and had invested billions of dollars in the startup. In return, Microsoft had secured access to OpenAI’s intellectual property and LLM models, which it distributed via APIs to Azure enterprise subscribers.⁹ A March 2023 contract with Microsoft JPMC ensured the protection needed to run client data on LLMs inside its own firewalls. “This was a big deal: it was the first time using confidential data in a scaled way on shared computing infrastructure. We needed a really trusted partner,” Waldron said. The decision to build JPMC’s own AI platform had a second benefit, Heitsenrether noted, “Ultimately, we’d like to be able to move pretty fluidly across models depending on the use cases to optimize capabilities, performance and cost.” Waldron added, “We knew that there were serious alternatives in the market, and we had no idea who was going to win. We used an aggregation strategy to abstract ourselves from the underlying providers.”

Appropriate Controls

As JPMC began receiving API access, teams from across the company began experimenting with GenAI. Waldron partnered with Heitsenrether and scoured the business units for signs of interest. Within weeks, a central intake portal had documented 1,000 ideas that businesses wanted to explore to develop into use cases. The firm also decided it needed to tailor its risk management approach for GenAI use cases. High-risk use cases, including those that were customer facing, went through the firm’s model review process, in which an independent risk management team validated the model.

Low-risk use cases, typically internally facing applications, required a risk and control assessment performed by the sponsor and the firm's control organization to confirm that the risks associated with the use case were appropriately controlled, typically by having a human in the loop. The firm's Compliance and Operational Risk organization reviewed all risk and control assessments. Bacon explained: "At JPMorgan, you cannot use an AI model without registering the model in our control process. Developers must adequately test their use case, and all use cases must go through a control process appropriate to the level of risk."

From a governance perspective, however, GenAI models presented unique challenges. Unlike traditional machine learning (ML) models that enabled *prediction* and *pattern recognition* on specific and narrow use cases, GenAI offered users *content creation* capability in the form of text, image, sound, and video generation across a variety of use cases.¹⁰ For example, the text generation function could be used to write an email, a memo, or to generate code. But because LLMs accommodated many inputs and delivered probabilistic output, repeated LLM completions could vary. Early releases of the models also delivered factually incorrect output, as they did not have live Internet access or were prone to hallucinate by making up text that otherwise appeared correct and convincing. Said Waldron, "We would not permit a solution to go live until this risk was appropriately managed commensurate with the level of human oversight. Especially for any client-facing solutions, understanding and managing this risk is of particular importance." Finally, risk teams worried about potential GenAI bias, and more generally about claims that a human in the GenAI loop would remove bias and error. "We know that our employees learn to trust quickly what their tools can do. That can result in not triangulating them with other sources," Bacon said.¹¹

AWM Chief Data Analytics Officer and Chief Information Officer Mike Urciuoli worked intensively on these topics for many months. He explained:

Within AWM, we immediately saw the extraordinary potential GenAI had for our business. Advice for our clients is rooted in language, and LLMs are a transformative technology in that domain. We immediately felt that it would bring the businesses that moved fast and innovated first immense competitive advantage, and we marshaled 150 data scientists and ML experts to lean into leveraging it.

But the technology was brand new, and the firm, the industry, and the regulators were all grappling with the challenges of hallucination and the necessary legal and control processes to address it. My team and I spent many hours creating a 12-level set of "guardrails," unique at the time, that reduced hallucination risk and improved answer quality. We were experimenting and iterating, learning from each other. There were no books to read. This was all new. GenAI models are incredible and broad, and these guardrails made something unpredictable—how LLMs answer questions—not only predictable, but safer to use.

AWM also mitigated risk by requiring humans in the loop. PB advisors were aware that they remained responsible for the guidance they offered clients, whether it included ideas generated with the help of GenAI or not. They were able to quickly access the citations and content GenAI responses were based on for more detail. Subject matter experts also reviewed a sub-set of GenAI responses, supplemented by systematic sampling as volumes grew. Among other things, journaling, logging, and audit features were built into GenAI processes for traceability.

Rewarding Inspiration

Despite many areas of active engagement across JPMC, skepticism and competing priorities slowed GenAI adoption in others. Some viewed GenAI as a reprise of the hype of the e-commerce dot-com era or that around blockchain technology.¹² Technologists and engineers involved in JPMC information technology (IT) systems focused on existing projects. There were also people who feared that the use of GenAI could eventually put their jobs at risk. Erdoes, however, was a quick convert to GenAI and spearheaded its adoption at AWM. With a laugh, she recounted, “I put the ChatGPT icon on the first page of my phone screen and moved Google and Safari to its second screen. My family thought I was nuts, talking to that thing all the time.” She also pushed to make AWM a first mover in the technology. Though some suggested that managing directors alone receive GenAI access, Erdoes was adamant from the beginning, Heitsenrether related, “Every single person—from administrative assistants to code writers—would have it, and AWM will be first.” Erdoes elaborated:

The processes in a larger organization are complex. New people come here every year and can’t believe we operate the ways we do. There are lots of good reasons for the roadblocks, procedures, and review committees. But, if we could do it faster, more creatively—that’s a whole new game. How are we going to play it? You need a whole program to get over the denial of “This can’t be right,” “It’s not as good as a human,” “It doesn’t test perfectly.” Have you met the regular humans doing the work today? They can’t always do it either. We can’t hit perfection every day, but we believe that the competitors who do this first will have a distinct advantage. Productivity gains are massive, and the models are moving so quickly that there’s danger in being a follower. If it becomes generative and self-learning, you have 200,000 people’s prompts and feedback helping extend our lead over the competition every day.

Heitsenrether enabled LLM access to groups throughout JPMC seeking greater access to GenAI. “[AWM] was the first division to use generative AI, piloting a generative AI ‘copilot’ for its private bank this summer [2024],” she said. “When the investment bank found out, they said, ‘Well, wait a minute, we want to be on there too.’ It does create a flywheel effect.”¹³

To ignite engagement, Erdoes and Urciuoli organized AWM “AI Days” in June and September 2023. For the June session, senior managers were instructed to have ideas to present and that participation would factor into their performance review. “It had buzzers, excitement, and early career professionals served as judges,” Erdoes said. Cross-functional teams—from the U.S., India, the U.K.—presented 30 demos, primarily featuring sales and portfolio management ideas. “The active groups had lots to show, but others showcased early ideas,” Urciuoli recalled. “That moment created a lot more inspiration but also pressure on all the teams to do something with the technology for fear of missing out, and it drove significant momentum in the business.” The second AI Day saw 100 more evenly representative use cases, many with prototypes well underway. Urciuoli said, “Any collaborative event that takes users and technologists out of their day-to-day jobs creates incredible things. In the early days of digital adoption, hackathons were tremendously productive, so we did the same for GenAI. Inspiration is contagious. The sharing inspires not only innovation but also user adoption in their own areas of activity.”

In addition to AI Days, a global *Shark Tank*-style competition for best GenAI application followed.^b AWM operating committee members evaluated the top six of 500 GenAI proposals submitted by early career analysts. Competition rules obligated JPMC to develop and execute the winning idea. “Some

^b *Shark Tank* was a reality television show featuring entrepreneurs pitching for funding and mentorship from a panel of investors.

contestants are just beginning their careers, and this shows that their ideas matter,” said PB Head of Digital Karen Donnelly. Other encouragement included training sessions and frequent communications, such as featuring users’ GenAI innovations in group emails or during morning meetings. An AWM internal website cataloged the most common LLM prompts. “We used peer competition as an adoption mechanism, too,” Urcioli said. “We showed statistics on who was using the technology, how often, and by group.”

Working for Erdoes, Donnelly was driven to make her team more effective. A PB advisor who had been promoted due to her repeated technological initiatives to remove advisor pain points, Donnelly now embraced GenAI. “If someone asks me to explain my role, I tell them that our team exists to make the advisor experience better,” she said. “They want smart reasons to call clients. They want to cover more clients to have more revenue claim. GenAI lets us do more business, better business, more things faster, with more leverage.” She fostered others with similar commitment wherever she found them. “When I took the job, someone named Mohamed Khalifa kept sending me ideas,” she said. “Who was this guy? He was levels down. I loved his enthusiasm. Now he reports to me and has leadership across the bank. Curiosity and a democratizing mindset are critical. That’s what we value. If you have a sleepy mindset, you will have sleepy results in the future.”

Bootstrapping an Operating Approach to GenAI

By mid-2023, through decentralized experimentation across JPMC, team leaders had gained enough insight into GenAI and its usage to agree on principles that would guide further development. Central among them was that broad access to GenAI would spur adoption and catalyze innovations that lifted productivity.

Democratized access By creating central access to LLMs through JPMC’s internally developed LLM Suite, the bank successfully sidestepped a potential stumbling block to adoption: fixed vendor subscription charges that were incurred regardless of usage. This avoided unproductive internal debate and gave JPMC teams access to GenAI for the negligible cost of the computing power they consumed using it.

Crawl, walk, run The practice of piloting and experimenting with GenAI initiatives created early feedback loops that identified and resolved obstacles before they reached scale. Dedicated working sessions gave cross-functional teams the opportunity to experiment and hash out design, testing, and implementation of GenAI applications without detracting from the team bandwidth needed to manage day-to-day bank operations or unexpected events, such as the collapse of three U.S. regional banks in spring 2023. Incrementalism rather than a rush to maximize impact proved helpful in negotiating governance challenges, too. Urcioli explored GenAI capabilities by ingesting non-sensitive JPMC documents rather than by seeking insights from more sensitive data sources.

Abstracting from use cases The proliferation of use cases revealed that most could be distilled to patterns that recurred across JPMC, such as information summary or techniques of GenAI dialoging. This led JPMC to begin constructing a platform that would avoid fragmentation by addressing use case patterns rather than their variety, in the process enhancing the scalability of AI applications. The result became LLM Suite, an integrated GenAI desktop assistant. In parallel, by the second AI Day, Urcioli said, “We found significant commonality across many projects being showcased. As a result, we

distilled the projects into a common set of patterns that every area could reference and asked teams to leverage these to build GenAI solutions^c for their specific businesses.”

Self-service The self-service model underlying LLM Suite and other JPMC GenAI applications allowed business users to experiment and solve their problems without needing data science expertise or multi-stage development and approvals. Having what Waldron called an “AI-first tool at scale on people’s desktops” significantly increased AI adoption and practice. With 200,000 users firmwide, JPMC data scientists could build plug-ins and distribute them at scale right away. LLM Suite was now a landing pad for solutions developed either centrally, by the business units, or by individual users.

Implementing GenAI

The Role of Information Technology

Speaking in late 2024, Beer said, “We invest \$17 billion in technology today. That was \$9 billion when I joined J. P. Morgan in 2014. We have integrated and substantially modernized our IT and data infrastructure, including a full private cloud and public cloud. We built a flexible, common tool chain for software development that could run at speed for the full complexity of regulation, technology, and cybersecurity controls we adhere to across the firm.”

Much of the impetus for these updates came from JPMC operating committee support that Beer elicited after sharing the commercial value that IT innovation occurring outside the firm could bring if applied to JPMC operations. “We laid the data on the table and showed the need for significant investment in an ongoing way, given the speed of evolution in technology today. That made inserting AI tools for engineers a lot easier,” she said.

By early 2025, JPMC technology had 63,000 employees focused both on operations—core infrastructure and software—and innovation and modernization efforts.¹⁴ Tasks that enabled GenAI were in the mix, including substantial sums to better organize and access JPMC data. Some of these efforts reflected ongoing commitments. In April 2024, for instance, CEO Jamie Dimon announced that public and private cloud providers would host 75% of JPMC data (up from 70%) and 70% of its applications (up from 50%) by year’s end.^d JPMC already employed over 2,000 AI, data science, and ML specialists, who were developing more than 400 GenAI applications.¹⁵ Since 2020, JPMC also had spent \$2 billion to build four private cloud data centers in the U.S., increasing its global footprint of data centers to 32 in 2024.¹⁶ The company managed nearly an exabyte of data, according to Beer.¹⁷

Laying a foundation for GenAI Reporting to Heitsenrether and Beer, Manoj Sindhvani was CIO of JPMC’s data and analytics office. He was focused on accelerating the development and impact of GenAI at JPMC by delivering an integrated data and AI platform that solved two problems. One was the challenge posed by the broad distribution of JPMC data in many on-premises, and public and private cloud-based systems—data that also could be difficult to access for processing. Another was proliferation of tools used by JPMC data scientists to build, test, and deploy traditional AI models and GenAI applications. Substantial progress on these goals had already been made, he said: “The platform allows data owners to publish the data and data scientists to request access to data, to train on it once

^c LLM solutions can include user-input driven and increasingly autonomous agents.

^d At the time, 80% of CCB production applications had migrated to the private and public cloud. Ninety percent of CCP analytical data was in the public cloud. JPMC, “Consumer & Community Banking: Investor Day 2024,” May 20, 2024, p. 2, <https://tinyurl.com/bdz94wu8>, accessed March 2025.

access is granted, and seamlessly deploy models to production.” A type of decentralized data architecture called a data mesh made JPMC data discoverable and accessible for training GenAI and traditional AI models, without the need to make copies. Governance was built in to link use cases to data and AI models.

The platform was being extended to accelerate build-out of agentic systems by leveraging foundational capabilities and registry of tools, agents, and prompts. “The platform is live, data are getting published, and the agentic builder is under development. All of this is happening at speed right now,” Sindhwani said. The platform had enabled a JPMC team to build in two days a new GenAI tool that gathered and synthesized information found on the Internet. “This would have taken much longer a few months ago,” he said. “We measure impact as time reduced from ideation to production. We expect that the development time to build new agents will be reduced by up to 50%.”

Benefiting from GenAI Beer also monitored four areas where GenAI was enhancing JPM’s technology functions: software engineering, infrastructure and operations, cybersecurity and controls, and employee experiences.

Software engineering LLMs assisted JPMC engineers with coding, code reviews, code translation (from one programming language to another), and quality testing. “Thirty percent of a software engineer’s time is spent coding, and 30,000 of our software engineers are using GenAI coding assistance right now,” Beer said. “Most firms have not scaled like this so far.” Moreover, she disclosed, 10% to 20% productivity gains from GenAI coding assistance were already evident.¹⁸ As a result, the firm was extending GenAI coding assistance from 20% to 100% of JPMC software developers. Global Chief Technology Officer Sri Shivananda reported to Beer and was responsible for JPMC’s technology architecture, research, and strategy. He said, “There is so much excitement around Generative AI that we are leveraging inspired teams across the organization to experiment. When and where we see promise from these experiments, we expand its usage across the organization. We fan out to discover good ideas, and fan those back in to scale them to the firm.” Test generation was an instance where GenAI was not assisting engineers by suggesting or analyzing code, but instead by acting directly to create new code from requirements, design, and research inputs. “Applying agentic architecture patterns is helping us enrich the software supply chain with builder/verifier approaches and workflow efficiency and effectiveness,” Shivananda said. “We are focusing a lot of effort into security, stability, and guardrails of our models. This will ensure that there are controls on how they interact with each other, with journaling, auditing, and accountability all built in. We must stand by this when we go external with it and the bar is raised; we must build trust and confidence mechanisms into AI.”

Infrastructure and operations As part of its daily operations, JPMC created logs of millions of transactions. Anomalies inevitably occurred, for any number of possible reasons—a server application upgrade, or a change on a network configuration device. When incidents occurred, LLMs combed the logs and identified the top three most likely causes, augmenting the work of existing ML models.

Cybersecurity and controls GenAI applications helped promote compliance and mitigate threat vectors by summarizing evidence of intent that the system was not designed to serve, or drift from that intent in underlying ML and AI models. Examples could be as simple as applications not using the latest version in JPMC’s code library. “One GenAI use case I am really excited about is our cyberthreat model,” Beer said. She explained:

Usually, security engineers would apply the principles of cyber governance that we abide by to threat patterns observed in cyber architectural diagrams, pictures, and schematics. Now, a cyberthreat tool that ingests all of that material, matches it up against those standards, and highlights the threat. In the future, we see the breadth of threat

modeling going way up, completely redesigning how we do risk management. We are in production with a cyberthreat model that will run every time we change code. We have over 4 million changes a year, so this will help us quickly understand any failures, or preemptively detect a network component failure.

For instance, Urciuoli said, “We get lots of alerts into our surveillance systems and supervisory [staff] spend significant time dispositioning false positives. Machine learning and GenAI are helping to highlight the relevant signals. Other systems couldn’t do that. AI helps prioritize the most important things to look at to find the real problems.”

As a result, Beer concluded, “GenAI is forcing us to look end-to-end at the software development lifecycle. That carries business process re-engineering implications as well. GenAI shows us things that cause us to ask why certain policies exist. Because AI is forcing everybody to look at the problem differently, it carries a halo effect of streamlining bureaucracy.”

LLM Suite

JPMC created its digital GenAI platform, LLM Suite, for the broad JPMC workforce rather than software engineers and AI model builders. Heitsenrether described LLM Suite as an LLM natural language interface designed to integrate with JPMC’s upgraded data platform and controlled organizational knowledge base. This maintained data security and logged user prompts and LLM Suite responses for compliance purposes while democratizing access to GenAI. Unlike publicly available LLMs, LLM Suite was embedded in existing JPMC desktop software; it was not a separate program to load and shut down. Heitsenrether explained, “It looks like a ChatGPT window, but behind that, we link to our data and retain the option to swap in and out new models based on their strengths. Think of it as a fund of funds model that we designed from the start to access the best tool for each task. Sign-in and access is via existing internal controls. Users just do their work at speed.” She elaborated:

We have been very focused on traditional AI for a long time. It is used in risk and fraud modeling, consumer credit, and AWM has been using it extensively. The biggest departure from traditional ML and AI is that the scale of GenAI becomes quite significant. The hurdle rate for intensive data science projects until now has been quite high, but with GenAI on most desktops in the firm, you unleash practitioners and no longer rely upon the scarce talent of data scientists to generate value. The resulting ideas and creativity are amazing, precisely because they come directly from practitioners in the business.

Alpha testing and governance requirements limited user numbers from a handful in January 2024, to 100 in April and 1,000 in May, when approval came to go live. Users then doubled in five days. In AWM, Urciuoli recalled, “We rolled it out to 30,000 of our people over six weeks. In the first month, we saw 250,000 prompts. So many people in our business reached out to thank us and say how impactful it was in their jobs. Then we started to see a million prompts per month.” Erdoes commented, “We now have agents, and bots, and a library of prompts. But if we hadn’t taken a bottom-up approach, we would still be in early stages of overthinking. You have to get buy-in for GenAI, and it will make the lives of those who find a good hook for it better. We went on a mission to find and eliminate all ‘no-joy’ work. Take that away, and you unleash even more potential in people.”

Typical users engaged LLM Suite to search, generate, and summarize text, or for general queries regarding JPMC policies, financial market products and dynamics, or legal and regulatory specifics. Demand for ever more information brought credentialed users access to JPMC research and knowledge repositories, earnings transcripts, news sources, and collaborative planning pages. Over time, LLM Suite capabilities would grow to encompass PowerPoint generation, Excel data analysis modules, and

sophisticated document analysis. Agentic workflows were the next horizon, and Heitsenrether's office was focused on developing the right governance and guardrails to ensure a balance between the strategic needs of the enterprise and what could be done at its edges. "In a regulated financial organization, you can't have chatbots roaming everywhere," she noted. "We welcome a free range of ideas, but that also needs the right central controls to avoid duplicative tools and models built to varying standards. We think about the best ways to do that night and day."

LLM Suite access across JPMC rose to 200,000 people between May and November 2024 (see **Exhibit 4**), or to two-thirds of the workforce.^e "Remarkably," Waldron pointed out, "regardless of how many people were on the platform, we consistently saw the same ratios: 30% were active users from the start, 10% were super-users, and the rest were slow adopters. That's traditional S-curve adoption. Once at scale, the 30% active-user figure rose to 50%." He continued: "We identified and interviewed the super-users and involved them in train-the-trainer sessions. These are the evangelists—frontier thinkers who enjoy innovation and imagination. They are first to push themselves with the new tools, and they have figured out how to turn LLM Suite into an asset for their team. Now that development cost has moved to zero, you release innovation and scale."

GenAI in Private Banking

In 2024, nearly 4,000 advisors provided 80,000 PB clients with advice, products, and services in support of client goals.¹⁹ In this process, the advisors drew upon AWM financial, insurance, and other expertise while cultivating client trust. Successful advisors were skilled at recognizing client wants and needs and matching the best possible offering to them, even for those who might lack familiarity with financial products or experience with wealth.

Robert (Bob) Blanch was JPMC PB head of investments and advice and one of its leading client advisors. He described JPMC private bankers as being among the smartest in the industry, with one of the strongest, most comprehensive platforms supporting them. This included a development program rooted in mentorship and learning by doing. He said:

Investment performance is quantitative, and clients will vote with their dollars. Our focus is more qualitative. Advice can be commoditized, but delivery is the differentiator. You have to know how to frame the questions, and there's a dynamic in pairing people on some client relationships. We spend a lot of time training our advisors in the J. P. Morgan way. We have an apprenticeship model, like a teaching hospital. A new advisor is paired with someone more tenured and follows a multiyear curriculum. There's a high correlation between those who are successful and those who go through the program.

Delivering good advice well involved knowing how available products mapped to client objectives, and what changing events implied for each of an advisor's several dozen clients. Erdoes explained:

When there is a big move in the stock market, do you know how hard it is to call people and tell them to sell it, GRAT it, collar it, or put it in a DAF trust?^f I do that repeatedly. You don't want to call the same person you spoke to two days ago. You don't want to call someone who doesn't make decisions; you want to focus on those who will act today. AI

^e Separate GenAI systems, supported by 20,000 software engineers, catered to the specific needs of JPMC's customer call centers, primarily to help their workers locate accurate, relevant information more quickly. Alexander Saeedy, "The Rise of Artificial Intelligence at JPMorgan," *The Wall Street Journal*, February 24, 2025, <https://tinyurl.com/4zt4vrca>, accessed March 2025.

^f GRAT referred to grantor retained annuity trusts, which helped minimize taxes on financial gifts to family members; DAF referred to donor-advised funds for charitable giving.

tools can tell you the 10 most useful calls to make to deliver daily personalized financial advice that's proactive. Advisors can tend to focus on their most interactive clients, and they need to have something to say to those they haven't spoken with in a while. These tools allow just that—on steroids.

This was the commercial imperative behind the advisor's role. Clayton Erwin was global head of investment sales for PB. He explained:

Private bank revenues come from deposits, loans, and investments. Advisors serve as our distribution team. Our goal is to deliver the products our clients need, create supporting collateral, and provide it to clients efficiently. Bob Blanch shouldn't have to rely solely on memory that a client said this, and it led to that. GenAI has a role to play in parsing client information intelligently to help allow them to match their needs to our products and let them execute directly. But will I as a client be comfortable receiving recommendations from a bot?

Private banking was on a path to discovering that answer. GenAI allowed clients to receive PB content at scale. Automated content personalization still lay in the future, but nothing stopped advisors from tagging clients interested in fixed income to receive articles from a JPMC expert on bonds or earnings call summaries. At the same time, advisors were able to capitalize on the nuance of personal relationships, whereas automated marketing could not. "Being able to customize marketing materials and presentations to suit client needs is a worthy objective," one advisor said. "We should ensure that the information still feels customized, and we need to be proactive in figuring out what piece of advice goes to what client on what day, recognizing that they have only so much bandwidth to absorb information from us."

For the moment, most GenAI applications available to PB advisors prioritized enhancing their client relationships over personalization strategies or client self-service (see **Exhibit 5**).

Technical Foundations for GenAI in the Private Bank

The release of ChatGPT was both exciting and timely for the 4,000 AWM technologists Urciuoli led. In the preceding years, AWM had focused on modernizing its core infrastructure, a crucial task that was now nearly complete. For instance, transitioning to the public cloud provided AWM scale and flexibility that significantly enhanced data processing capabilities and made AI deployment more efficient. Data modernization ensured that the data infrastructure was robust and able to support advanced AI applications, including GenAI. In addition, the microservices architecture of AWM's new, state-of-the-art Connect platform for the Private Bank allowed seamless integration with GenAI APIs and flexible deployment of AI solutions at scale.

AWM Technology also had laid important groundwork through prior investments in talent, including over 100 data scientists and an equal number of AI/ML experts. "Key leads in AI had been exploring and building solutions with LLMs years before the release of ChatGPT," said Vrinda Menon, a chief technology officer whose team had built some early solutions. This early exploration, combined with existing expertise and an innovation mindset, provided a head start in integrating GenAI.

Coach

Coach was more ambitious than some of those applications. In early 2025, it was one of the largest GenAI applications at JPMC and had taken multiple AWM teams—each combining technologists, data

scientists, product managers, and subject experts—over a year to create. Coach was embedded in Connect, the digital platform for PB advisors. Connect allowed advisors to perform multiple tasks, from setting up new accounts and structuring new investments to enabling lending, banking and customer account management. Coach combined Connect’s client profiles, account data, and summaries of client-advisor interactions with GenAI tools and JPMC wealth management information, supplemented as needed with published sources. Launched to 3,000 client advisors in October 2024, its immediate aim was to automate and scale delivery of accurate, complete, easy-to-access information.

This solved several problems. Advisors often had to obtain answers for client needs either by calling JPMC subject matter experts or by manually searching content libraries. Delays in receiving needed information could hinder client service. Rapid access to JPMC-vetted resources saved advisors time better spent deepening existing client relationships and developing new ones. Indeed, AWM leadership hoped that Coach and other GenAI tools would create efficiency gains sufficient for advisors to grow their client rosters by 50%. Importantly, Coach could make fuller use of JPMC client data to help understand clients and prospects, too.

Coach responded to advisor queries by tapping any of more than two dozen agents. For example, if an advisor asked how best to prepare for a client discussion about a particular sector, Coach would first parse the advisor’s query through a natural language interface. Its orchestrator agent then would send the relevant queries to internal agents able to address them, drawing on both structured and unstructured data. Coach would then synthesize the agent responses and send its response back to the natural language interface (see **Exhibit 6**).⁸

Building GenAI Solutions

In 2020, AWM had introduced the advisor chatbot Casey, developed by Menon and her team, to provide quick access to a fixed set of capabilities, such as checking a client’s account status or generating specific documents. Said Urciuoli:

Advisors loved it and it was used heavily. While it had a limited set of capabilities, it confirmed the paradigm of a chat window providing easy access. Many of us imagined GenAI creating the next generation of Casey. When GenAI came out, we saw it would allow us to do much more than Casey did, with much more flexibility, while preserving the chat window interface. Making that mental connection was part of our “a-ha” moment with GenAI, because we saw that it would empower advisors to go faster by putting information at their fingertips, giving clients a better experience with better information. GenAI would be Casey 2.0.

Early in 2023, AWM technology teams received access to OpenAI’s LLM through Microsoft APIs and began exploring use cases and building applications that harnessed AWM investment and other expertise. Among the earliest GenAI projects approved for JPMC pilot was PlanningGPT, an AWM solution created to answer questions about estate planning using content gathered from PB subject matter experts. As part of its testing, 600 PB advisors peppered the solution with 4,300 questions relating to 350 documents that it had ingested. Subject experts then reviewed the accuracy and completeness of 3,700 responses, selecting both from responses considered potentially higher-risk than average, and from a random sample of all queries. Analysis found 50% of PlanningGPT responses to be correct and fully robust. Another 40% of responses were judged correct and partially robust. Eventually, PlanningGPT ingested 90% of JPMC’s available estate planning texts, each selected and

⁸ Applications could include one or more agents and referred broadly to any software using GenAI techniques. Orchestrators played a coordinating function among multiple agents, LLMs, and other GenAI technologies.

validated by subject experts from documents approved for external distribution. Turbo, an improved version of PlanningGPT, supplemented JPMC content with public information, boosting the share of LLM responses deemed fully robust from 50% to 90%. Partially robust answers comprised another 6%.

The pilot set several notable precedents for all later efforts. First, it created standards for subject expert involvement in reviewing LLM responses across AWM. Second, the exercise helped generate a feedback mechanism that aided LLM training. Third, it prompted development of a content upload tool to ingest new and remove expired content. Fourth, it led to systematic ingestion of external content.

Perhaps most importantly, subject expert teams embedded with technologists retained responsibility for the GenAI solution's creation, testing, and for the content selected and ingested—because they knew that content best. Urciuoli had pushed to decentralize AWM data scientists and AI/ML specialists to subject expert teams, recognizing that this federated approach would allow development of solutions for different subject areas in parallel, accelerating overall progress while equipping the teams with important capabilities. He said, “That empowered people who work for me. Technology and data science are powerful tools. You want to give people the technology and inspiration and let them run with it. That’s the scalable approach. That’s bringing them along for the journey. That’s what we want to do in every area.”

Other AWM teams developed GPT solutions too, each treating a different topic—investments, philanthropy, taxes, and a recap of JPMC’s daily market views—in isolation. Aside from the need to chunk JPMC content for LLM ingestion, the cost to create each was low. The proliferation of chatbot windows confused advisors, however, who did not always realize which tool to use when.

In summer 2024, a PB technology team created the Coach framework to put the individual solutions to collective use behind a single chatbot interface. This involved creating a coordinating orchestrator to interpret advisor queries and route them to the relevant solution for processing. Because the power of LLMs rested in part on adapting to the queries and prompts they received, Coach benefited not only from access to an advisor’s previous workflows, but also from advisor input. Onscreen buttons, such as “Ask Coach” and “Always learn more about this topic,” were added to elicit interaction. Coach could also query advisors (e.g., “Do you know you have three clients with birthdays today?”) and propose actions with commercial benefit.

Testing

Over time, the teams codified testing to include all questions from pilot users and Coach answers. The pool of participants expanded to include technology experts and low-technology users in addition to subject matter experts. The number of queries evaluated rose to multiple tens of thousands, all from subject experts. Incomplete answers were escalated for improvement. Blatantly incorrect responses were obvious, Donnelly related, and they typically involved the omission of critical detail. Often, the cause lay in outdated source documents, which led to improved accuracy in materials received from marketing and other content partners. Unexplained wrong answers were examined at length due to compliance, risk, and control requirements.

Usage

Advisors used Coach primarily to generate memos of client calls, summary records of advisor/client discussions, and a list of follow-up tasks. Advisors could review and edit each memo before saving it, often sharing a copy with the client. Data crawls were the second most popular use of Coach. For all JPMC text Coach returned, hyperlinked citations and page references to the source were the norm. Every Coach response distinguished between JPMC and any supplemental public sources cited, adding a disclaimer for the latter (see **Exhibit 7**). Investment information comprised 85% of crawl

queries. Natural language queries were third-most common. These might include requests for a list of clients the advisor had not spoken to recently or for those underexposed to an asset class.

More broadly, Coach helped advisors in three ways: (a) locating and providing links to relevant JPMC externally approved resources that could be shared with PB clients; (b) educating advisors, using both internal and external resources, on content, policy, and current events relevant to their clients; and (c) generating knowledge from JPMC and external sources that was personalized to client needs. As a result, PB Head of AI and Innovation Mohamed Khalifa said, “The internal cost of expertise for basic information is going to zero because advisors no longer need to disturb or consult experts. Advisors can get everything they need from Coach themselves, personalized for their clients.”

Extending the Frontier

Erdoes now proposed Coach as a vehicle to deploy GenAI much more deeply in support of the client experience. Efforts were underway to hyper-personalize Coach by drawing upon advisor notes, client transaction patterns, and appropriate client-specific data. This would provide all AWM advisors the degree of client awareness that its best performers were known for, to provide information and guidance tailored to individual client needs and temperament, particularly in response to market developments.

A potential next step could be making Coach’s existing capabilities directly available to clients to give them access to the same product and policy data, resources, and empowerment available to advisors. In some respects, sharing with clients the benefits of GenAI as its capabilities evolved from assisting users to acting on their behalf was a logical next step (see **Exhibit 8**). Heitsenrether was not unsympathetic. She observed: “AI does give you the ability to scale and provide more personalization. Most of us, I think, would prefer if you can just go online to a digital channel and get the answer fast. The AI can help to discern your question and get you an answer in a digital channel more quickly.”²⁰ However, she also emphasized that the technology was not yet ready to be client-facing.

Blanch noted, moreover, “We did not embrace self-service for a very long time.” “There was a belief that clients didn’t need online trading.” Donnelly added, “We didn’t offer clients full online trading until March 2024. That was a problem. Since that time, we have executed 120,000 client-initiated trades. Many clients want more self-service, whether that means sending wires or changing their address.”

Technology solutions would continue to be implemented in close partnership with the business, Beer noted: “We want to democratize implementation, but with the guardrail of a platform.”

Others believed that access to Coach content would help engage rather than decouple client-advisor relationships. Donnelly said, “I don’t think clients will use the tool instead of advisors. The technology simply removes basic information issues and allows solid interaction. Clients could ask for our view on gold and could choose to buy gold online. The advisor would get an alert of the transaction. That would keep advisors up to date on clients’ asset allocations.”

A potential step toward GenAI-assisted advising might be to introduce a voice-enabled version of Coach into the client-advisor relationship. In this scenario, GenAI might serve the relationship by supplying information as part of client discussions with advisors as needed. But whether voice-enabled GenAI was a preface to client self-service or to creating more effective advisors remained untested.

The effects of GenAI on workforce size remained an unknown, as were their implications for the apprenticeship model common to PB advisors and other JPMC jobs. In June 2024, Citigroup estimated that GenAI could displace 54% of banking roles, augmenting those roles in just 12% of cases.²¹ The process re-engineering that Beer had observed as CIO was now occurring among AWM subject expert teams, whose lengthy analytical reports and position papers might no longer be superior to more

structured tabular data as the best way for GenAI to parse, organize, and deliver JPMC knowledge. Erdoes acknowledged that hiring plans had been shifting, but she viewed the changes GenAI brought as increasing positive pressure on bank managers for good performance. “We have the benefit of growth here, which means, generally, that we are hiring,” she said. “So, if AI makes things better, it’s not like I will see 20% reduction in time spent and reduce 20% of people, it’s that I don’t have to hire more. Let AI eat your job; we have lots of other jobs here for you to do. Your job won’t be taken by AI; it will be taken by a person mastering the use of AI.”

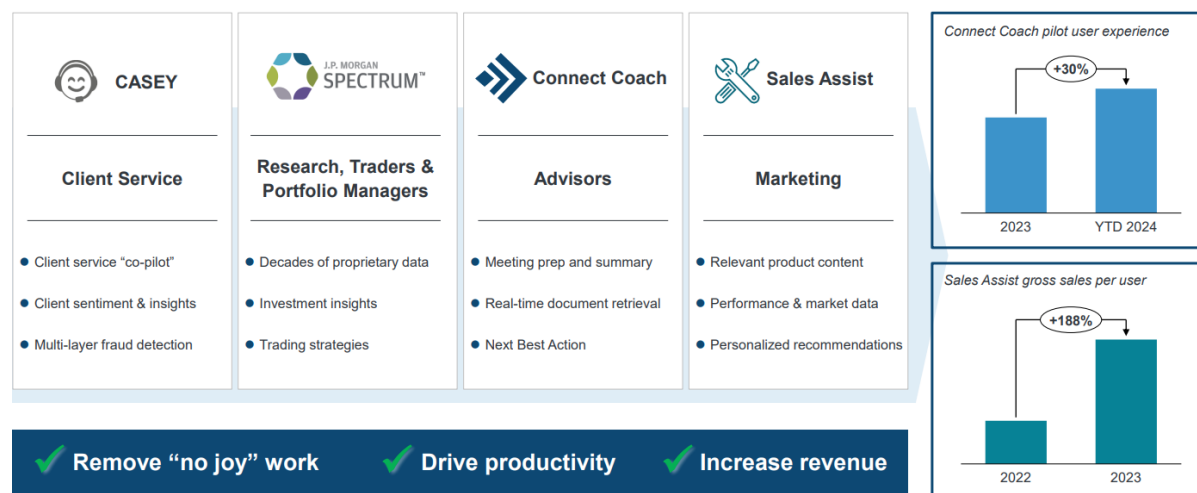
Dimon, too, was emphatic about GenAI and its impact. He said:

We are going to deploy this tool. It’s going to change over time. The cost may be more than we thought, it may be less than we thought. Again, I am not going to worry about it. We’re just going to be the best at it. It will affect every job out there; don’t put your head in the sand. A lot of companies don’t want to talk about it openly because, “We’re scaring people about jobs.” It is going to replace jobs. So did the tractor, and fertilizer. [. . .] The more I know about it, the more I can plan for it, let attrition be my friend, and where necessary, redeploy, retrain, et cetera. [. . .] We’re going to use it, and there won’t be one job that won’t be touched by it.²²

Exhibit 1 JPMC Operating Committee Members, March 2025

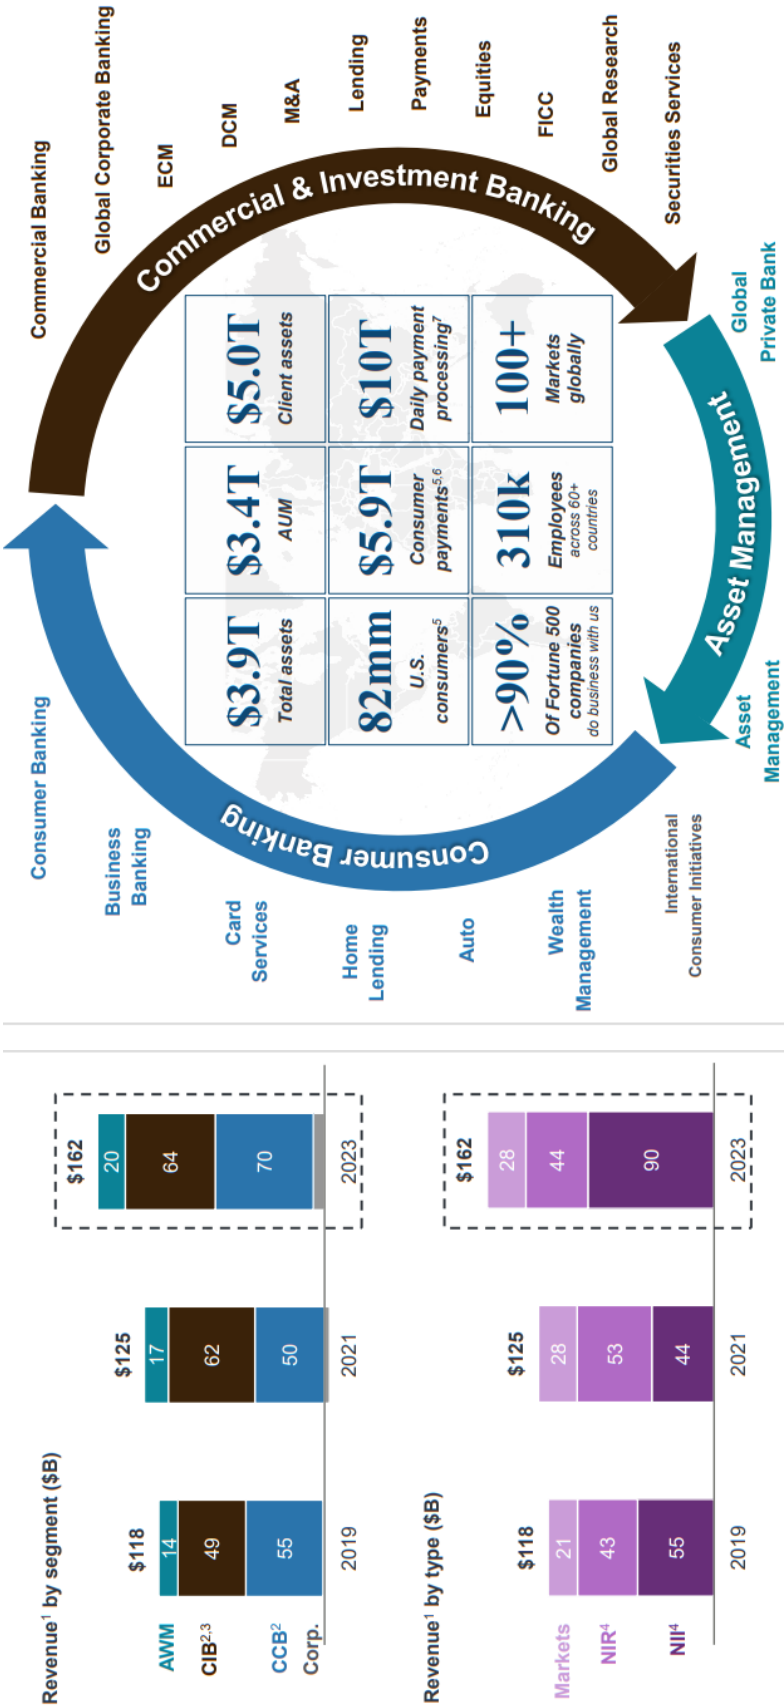
Name	Title	Joined JPM
Jamie Dimon	Chairman and CEO	2004
Ashley Bacon	Chief Risk Officer	1993
Jeremy Barnum	Chief Financial Officer	1994
Lori Beer	Global Chief Information Officer	2014
Tim Berry	Global Head of Corporate Responsibility and Chairman of the Mid-Atlantic Region	2017
Mary Callahan Erdoes	CEO, Asset & Wealth Management	1996
Stacy Friedman	General Counsel	2012
Teresa Heitsenrether	Chief Data and Analytics Officer	1987
Marianne Lake	CEO, Consumer & Community Banking	2000
Robin Leopold	Head of Human Resources	2010
Douglas B. Petno	Co-CEO, Commercial & Investment Bank	1990
Jennifer Piepszak	Chief Operating Officer	1994
Daniel Pinto	President and Chief Operating Officer	1983
Troy Rohrbaugh	Co-CEO, Commercial & Investment Bank	2005
Sanoke Viswanathan	CEO, International Consumer & Wealth	2010

Source: JPMC, "Our Leadership: Operating Committee," <https://tinyurl.com/4xd422c6>; LinkedIn.com profiles; "Marianne Lake, Jennifer Piepszak," Forbes.com, <https://tinyurl.com/bdt48m8t>; all accessed March 2025.

Exhibit 2 Selected AWM GenAI Applications, May 2024

Source: JPMC, "Asset & Wealth Management: Investor Day 2024," May 20, 2024, p. 14, <https://tinyurl.com/5dyvpfdj>, accessed January 2025.

Exhibit 3a Overview of JPMC Business Activities (in \$ billions and \$ trillions, various years)



Source: JPMC, "2024 Investor Day: Transcript," May 20, 2024, p. 4, <https://tinyurl.com/mtaxbbm3>, accessed January 2025.

Notes: AWM: Asset & Wealth Management; CIB: Commercial & Investment Bank; CCB: Consumer & Community Banking; Corp.: Corporate; NIR: Non-Interest Revenue; NII: Net Interest Income; AUM: Assets under Management; ECM: Equity Capital Markets; DCM: Debt Capital Markets; FICC: Fixed Income, Currency, and Commodities.

Exhibit 3b JPMC Summary Financial Data by Business Segment, 2022-2024 (in \$ billions and %)


	Consumer & Community Banking			Commercial & Investment Banking			Asset & Wealth Management		
	2024	2023	2022	2024	2023	2022	2024	2023	2022
Total net revenue	71.5	70.1	54.8 ^a	70.1	64.4	59.6 ^a	21.6	19.8	17.7
Total non-interest expense	38.0	34.8	31.2 ^a	35.4	34.0	32.1 ^a	14.4	12.7	11.8
Pre-provision profit/loss	33.5	35.3	23.6	34.8	30.4	27.6	7.2	7.0	5.9
Provision for credit losses	10.0	6.9	8.8	0.8	2.1	2.4	-0.1	0.2	0.1
Net income/loss	17.6	21.2	14.9 ^a	24.8	20.3	19.1 ^a	5.4	5.2	4.4
Return on equity	32%	38%	29%	18%	17%	14%	34%	31%	25%

	Corporate			Total		
	2024	2023	2022	2024	2023	2022
Total net revenue	17.4 ^b	8.0	0.1	180.6 ^b	162.4	132.3
Total non-interest expense	4.0 ^c	5.6	1.0	91.8 ^c	87.2	76.1
Pre-provision profit/loss	13.4	2.4	-1.0	88.8	75.2	56.1
Provision for credit losses	0.0	0.2	0.0	10.7	9.3	9.7
Net income/loss	10.6	2.8	-0.7	58.5	49.6	37.3
Return on equity	NM	NM	NM	18%	17%	14%

Source: Adapted from JPMC 2024 Form 10-K (filed February 14, 2025), p. 72, <https://tinyurl.com/yc55kr25>, accessed March 2025.

Notes: (a) In the first quarter of 2023, the allocations of revenue and expense to CCB associated with a Merchant Services revenue sharing agreement were discontinued and are now retained in Payments in CIB. Prior-period amounts have been revised to conform with the current presentation. (b) Included a \$7.9 billion net gain related to Visa shares recorded in the second quarter of 2024. (c) Included a \$1.0 billion contribution of Visa shares to the JPMorganChase Foundation recorded in the second quarter of 2024.

NM: not meaningful.

Exhibit 4 LLM Suite Users, January-November 2024

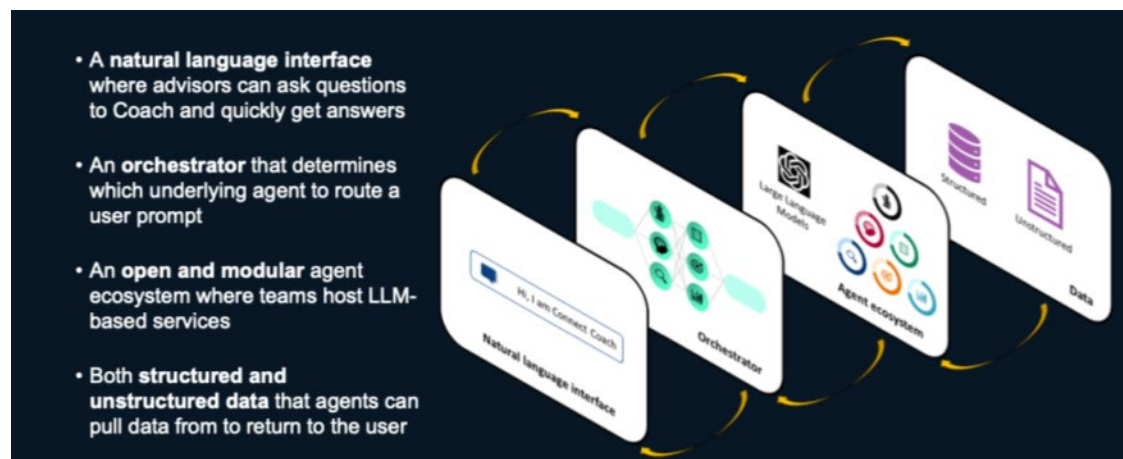
Source: Casewriters, based on company interviews.

Exhibit 5 Selected AWM GenAI Tools, 2024

- **Client360:** A customer relationship management (CRM) initiative that consolidates client interactions, aiming to optimize engagement by suggesting intelligent next actions to advisors. It integrates data across all client touchpoints, enhancing the client experience.
- **CoachAssist:** Categorizes client calls as “great,” “mediocre,” or “requiring improvement,” as a tool to improve client interactions and internal coaching.
- **Private Bank IQ:** Highlights top opportunities for advisors in response to events as diverse as a mortgage rate change or a localized weather event. Integrated with the email system so that the sales messages can be suggested to advisors and sent to selected clients with a mouse click.
- **SalesAssist:** Automates call transcription in real time and conveys audio suggestions to client advisors during client calls. SalesAssist allows advisors to remain present with clients by reducing the need for and distraction of searching for information in mid-conversation.
- **Sales Copilot:** A planned CRM tool to optimize the efficiency and effectiveness of client sales visits and interactions by advising grouping on travel and meeting arrangements by client location and availability.
- **SmartBrief:** Automates routine preparatory tasks such as compiling agendas and meeting notes, allowing junior staff more time for client engagement in the field with senior bankers.

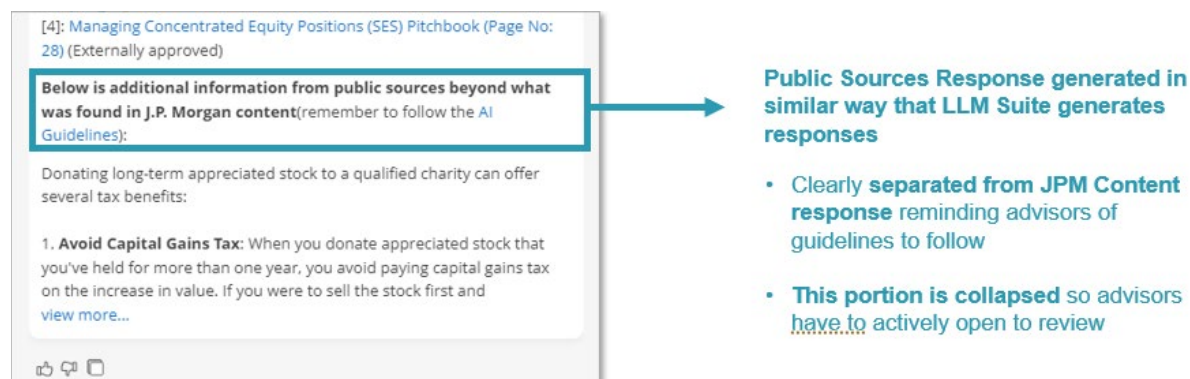
Source: Casewriters, based on company interviews.

Exhibit 6 Schematic Description of Coach Operations




Source: Company documents.

Exhibit 7 Treatment of Public Source Data in Coach



Source: Company documents.

Exhibit 8 Potential Path of GenAI Usage for Businesses



Source: Casewriters, based on company interviews.

Endnotes

- ¹ JPMC, “Chairman and CEO Letter to Shareholders,” April 8, 2024, <https://tinyurl.com/yuuvc9r3>, accessed January 2025.
- ² JPMorgan Chase & Co., “2024 Investor Day: Transcript,” May 20, 2024, p. 4, <https://tinyurl.com/mtaxbbm3>, accessed January 2025.
- ³ JPMC, “2024 Investor Day: Firm Overview,” May 20, 2024, slides 4, 18, <https://tinyurl.com/2s8ba6f7>; “League Tables,” *Financial Times*, <https://tinyurl.com/2kdj45wu>; Nathan Reiff, “10 Biggest Banks in the World,” Investopedia.com, October 13, 2024, <https://tinyurl.com/ysvybe5u>; all accessed January 2025.
- ⁴ Evident, “The Evident AI Index,” October 2024, <https://tinyurl.com/rn27enhk>, accessed January 2025.
- ⁵ Matt Ashare, “JPMorgan Chase Leads Banking Sector in AI Adoption: Report,” Banking Dive, October 18, 2024, <https://tinyurl.com/bde2dcwm>, accessed January 2025.
- ⁶ JPMC, 2024 Form 10-K (filed February 14, 2025), pp. 56, 84, <https://tinyurl.com/yc55kr25>, accessed March 2025.
- ⁷ JPMC, 2024 Form 10-K, p. 77.
- ⁸ JPMC, 2024 Form 10-K, p. 73.
- ⁹ Jordan Novet, “Microsoft’s \$13 Billion Bet on OpenAI Carries Huge Potential along with Plenty of Uncertainty,” CNBC.com, April 9, 2023, <https://tinyurl.com/5ceymacc>, accessed January 2025.
- ¹⁰ Hugh Son, “JPMorgan Chase Is Giving Its Employees an AI Assistant Powered by ChatGPT Maker OpenAI,” CNBC.com, August 9, 2024, <https://tinyurl.com/srpvk5t4>, accessed January 2025.
- ¹¹ Fabrizio Dell’Acqua, “Falling Asleep at the Wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters,” Harvard Business School Laboratory for Information Science working paper, 2022, <https://tinyurl.com/cdzpx57e>, accessed March 2025.
- ¹² Adam Hayes, “Dotcom Bubble Definition,” Investopedia, May 31, 2024, <https://tinyurl.com/2p8ur2zb>; Grant Gross, “RIP (Finally) to the Blockchain Hype,” CIO.com, March 6, 2025, <https://tinyurl.com/4nr5ayws>; both accessed March 2025.
- ¹³ Michelle Abrego and Bianca Chan, “JPMorgan’s AI Rollout: Jamie Dimon’s a ‘Tremendous’ User and It’s Caused Some ‘Healthy Competition’ among Teams,” Business Insider, November 22, 2024, <https://tinyurl.com/mskmac6f>, accessed March 2025.
- ¹⁴ Yamini Kalra, “Case Study: JP Morgan’s \$17 Billion Tech Push,” CIO.inc, November 18, 2024, <https://tinyurl.com/y6zuupb7>, accessed March 2025.
- ¹⁵ Matt Ashare, “JPMorgan Chase CEO Wants More Cloud to Fuel AI, Analytics,” CIO Dive, April 12, 2024, <https://tinyurl.com/5n6th6p7>, accessed January 2025.
- ¹⁶ Ashare, “JPMorgan Chase CEO Wants More Cloud to Fuel AI, Analytics.”
- ¹⁷ JPMorgan Chase, “JPMorganChase Global CIO Lori Beer Showcases Strategic Advancements at AWS re:Invent,” December 5, 2024, <https://tinyurl.com/49d743mp>, accessed March 2025.
- ¹⁸ Haripriya Suresh, “JPMorgan Engineers’ Efficiency Jumps as Much as 20% from Using Code Assistance,” Reuters, March 14, 2025, <https://tinyurl.com/249w7bzs>, accessed March 2025.
- ¹⁹ Jerry Dubrowski and Pan Hazard, “JPMorganChase Opens Flagship Financial Centers for Affluent Clients,” Chase Media Center, October 16, 2024, <https://tinyurl.com/367jaajf>; “The World’s Best Private Bank,” Euromoney, March 22, 2024, <https://tinyurl.com/2d56db96>; both accessed March 2025.
- ²⁰ Alexander Saeedy, “The Rise of Artificial Intelligence at JPMorgan,” *The Wall Street Journal*, February 24, 2025, <https://tinyurl.com/4zt4vrca>, accessed March 2025.
- ²¹ Jack Kelly, “AI Could Displace More than 50% of Banking Jobs, According to New Citigroup Report,” Forbes.com, June 20, 2024, <https://tinyurl.com/3dahavrv>, accessed January 2025.
- ²² Business Today, “Jamie Dimon on AI Impact: It Will Replace Jobs in Banking and Financial Services,” YouTube, published March 15, 2025, <https://tinyurl.com/4j6cwx2>, accessed March 2025.