

AI, Labor, and Economy Case Studies In-Brief

Axis Bank



PARTNERSHIP ON AI

Axis Bank is the third-largest private-sector bank in India, with 20 million customers and signs of substantial growth in the near future. Axis Bank's implementation of a chatbot powered in part by artificial intelligence (AI) — part of a broader effort to improve customer service quality, costs, and scalability — offers a glimpse into the labor and economic implications, as well as organizational challenges, of a company building and deploying an AI solution.

Case background

Founded in 1993, Axis Bank has recently sought ways to introduce scalable and cost-efficient customer service solutions to support its rapid customer growth. Keeping its existing operating model - composed of branch agents, and third-party-managed email and call center services - would require the bank to grow its customer service team headcount to keep up with this growth. Axis Bank's email customer service also had room for improvement, with low satisfaction scores from customers and long response times. To address these challenges, Axis Bank undertook a broad set of customer service initiatives in late 2016 and 2017 (some not AI-driven, and some AI-driven) including a Manual Chat service,¹ a customer Self-Serve Q&A Platform,² improvements to the existing Interactive Voice Response (IVR)³ system at the call center, and an AI Chatbot implementation called Axis Aha!. The AI Chatbot was an integral part of this effort to address the growing volume of customer inquiries in a scalable way through automation. (In this case, when we refer to the impacts of AI at Axis Bank, we generally are referring to the Axis Aha! Chatbot, though IVR also contains functionalities enabled by AI, such as audio recognition.)

To build the AI Chatbot, the project team used several machine-learning algorithms, as well as rule-based systems, to respond to customers' inquiries. Natural language processing/natural language understanding and natural language generation were used to understand and generate responses to customers over chat in the 'natural' forms customers may type. Neural networks were used to extract information from unstructured email service transcripts, banking documents, and banking Q&A web pages in order to generate the AI Chatbot's training data. Rule-based techniques were used to fill in manual answers for banking questions that the chatbot had low confidence were correct. In determining where to focus, Axis Bank prioritized use cases and transactions based on inquiry volumes, expected user adoption, and implementation complexity. After launch, the AI Chatbot was updated in 15-day release cycles to launch new features and improve the system's performance after exposure to new training data.

This case study documents the process Axis Bank used to build the AI Chatbot and explores its impact on business results and workers, both inside and outside the bank.

¹ Manual Chat (launched December 2017) is an online messaging platform for select banking products where customers can instant message with live customer service agents.

² Self-Serve Q&A Platform (launched September/October 2017) allows customers on Axis Bank's support website to find informational answers to frequently asked questions such as how to request a checking account or block a debit card.

³ IVR, short for Interactive Voice Response, is an automated, voice-based system that allows customers to navigate within a call center customer service channel.

Key observations and findings:

- **Implementation challenges:** Axis Bank faced numerous challenges as it implemented the chatbot. These challenges ranged from technical development difficulties and regulatory considerations to consumer education and stakeholder management. Given the system's complexity, training the AI Chatbot required generating new training data and adopting staged release cycles every 15 days with manual reviews. This development rhythm was substantially different than originally anticipated. The complexity of the application domain not only required changes to implementation processes, but also extended the challenges and timeline of development. Axis' banking-specific vocabulary necessitated creating the right corpus of materials for the chatbot and attention to the regulatory issues encountered. Axis and their technology vendor responded to potential regulatory and compliance requirements and risks by creating multiple testbeds, bringing in expert third parties to review the code, and reviewing the AI's learnings and updated behaviors before promoting new functionality to the broader public. In response to regulatory and compliance considerations, including requirements to store data locally, the AI Chatbot had to be run on site, as opposed to having a cloud-based infrastructure, though the latter may have been preferable from a performance perspective. An additional aspect that complicated the implementation was that Axis Bank actively managed and addressed customer feedback on Axis Aha! in an early beta release. For instance, Axis realized that customers wanted to execute tasks with the chatbot, not just retrieve information. Accordingly, the development team was compelled to build new functionalities into the chatbot.
- **Business and productivity impact:** The business impact of the AI Chatbot alone is difficult to precisely measure, as it was rolled out as part of a broader set of customer service initiatives, including Manual Chat and Self-Serve Q&A Platform launches, as well as improvements to existing Interactive Voice Response (IVR) at the call center. This set of customer service initiatives (including the AI Chatbot) helped reduce customer service staffing costs, while revealing new customer insights. Third-party customer service costs were reduced by 26 percent for a year's time, driven by a net reduction in (third-party) customer service headcount contracted, despite 10 percent year-over-year growth in customer service volumes. Axis Bank reports it was able to handle its growing customer service volumes with fewer customer service agents, driven by a migration in volumes from human-enabled to automated customer service channels, implying an increase in the labor productivity of the overall customer service function.⁴ Migration of customer service volumes to automated channels also improved customer service quality, including time to serve customer service inquiries and the average response time. Through enhanced customer service availability and diversity of automated customer service channels, Axis Bank expects overall customer satisfaction to improve over time.
- **Workforce and labor impact:** Axis Bank's customer service initiatives (including the AI Chatbot) had far greater workforce impact outside the bank than inside it. As a result of the automation initiatives, the email customer service team at the third-party vendor was cut by about 190 full-time equivalents (FTEs),^{5,6} down 76 percent to roughly 60 FTEs from a total of roughly 250. Interviewees reported that a number of these employees were re-trained for other jobs at the vendor or were moved to other Axis customer service channels, although this was not externally verified. Coinciding with labor displacement driven by the automation initiatives, Axis Bank reported that the AI Chatbot deployment also created more than 40 temporary jobs, through contracting with third-party technology and advisory service firms. Such jobs might include software developers, data scientists, and technology experts who developed and tested the chatbot. Internally, an estimated six to eleven Axis Bank FTEs were also affected by the chatbot implementation, largely in the form of reassignments and minor retraining without immediate job loss.

⁴ We do not have a measure of hours worked to estimate the increase in labor productivity precisely.

⁵ FTE refers to full-time equivalent employees at the company. The term is a business acronym and is a conventional unit of measure to compare workloads across different business contexts. Because labor may be undertaken by part-time employees, it is useful to standardize work amounts across full-time equivalents, rather than total worker counts.

⁶ These agents were redeployed to other Axis customer service channels or to other accounts at the third-party vendor.

Conclusion and lessons learned

Axis Bank developed an AI-enabled chatbot to serve customer support inquiries in a scalable way. The perceived need for an AI Chatbot reflects broader industry trends within the expanding Indian economy, such as growing consumer digital engagement. The diffusion of AI technologies in the banking sector aligns with the Indian government's vision of a "Digital India," in which AI can help people to overcome barriers to financial access.

As noted, Axis Bank's implementation of the bot was not without its challenges and had notable impacts on employment that rippled beyond the firm. As the case study reveals, training the chatbot for this specific industry context required significant time and resources. In general, building chatbots to execute business functions typically requires large amounts of customization and dataset generation; Axis was no exception.

Realistic expectations and real results: According to Axis Bank, in order for AI chatbots to work effectively, organizations must have a broader strategy, and define the precise role expected of the chatbot. Within Axis Bank, the AI Chatbot is part of a broader set of customer service automation initiatives, and thus is not expected to fulfill all possible functions. Still, and in conjunction with these other initiatives, the AI Chatbot implementation drove productivity gains at Axis Bank, by reducing service costs and elevating total customer service queries. Though productivity gains of AI applications and automation initiatives may not yet be visible in the broader economy, Axis Bank's case raises the possibility that productivity gains from automation suites, including AI, will spread in the future with increased diffusion of such technologies to other sectors.

Mixed labor implications — expanding beyond the boundaries of the firm: The AI Chatbot did not have major internal employment implications for Axis Bank. However, the labor implications external to Axis Bank yet still within the firm's business ecosystem were considerable. The case of Axis Bank raises questions about the future of the workforce, particularly for customer service. Axis Bank's use of automation and AI had far greater impact externally than internally on job losses, job gains, and reassignments. Only six to eleven internal employees saw their jobs change within the bank, and none experienced full job-loss. Accordingly, the bank experienced fewer change-management challenges than it may have otherwise, given the outsourced nature of this labor. For the Indian economy as a whole, with a large working demographic engaged in both in-house and outsourced roles, the long-term impacts of widely diffused AI use remain to be seen.

With regard to technical expertise, building an AI Chatbot did not require Axis Bank to invest in hiring AI experts into the bank. Instead, Axis partnered with external technology vendors to develop and test the new AI Chatbot. While chatbot development did demand specialized labor in the short-term, the ongoing need for this labor may decrease once these specific products and services move into production. For instance, interviewees suggested that the analysis, testing, and development of AI Chatbot iterations are being increasingly automated. Eventually, improvements to Axis Aha! may enable it to address more complex customer service tasks and ultimately replace Manual Chat agents. If the Axis Aha! Chatbot becomes able to converse in multiple languages, it may also potentially replace the call center in the long term. In this sense, Axis may not yet have experienced the full extent of Axis Aha!'s workforce ramifications. Still, offering the AI Chatbot in multiple other languages, such as in Hindi, would take significantly more development time. One implication, as an Axis executive put it, is that call centers "may never go away in the Indian context."

Open questions and future research

This case raises a number of unresolved questions for further research. Some of these questions pertain specifically to the labor experiences and implications for workers on the ground: How severely will front-line customer service employees (for instance, in-branch support specialists) be impacted in the medium to long term as deployment of AI chatbots becomes more widespread? How will the broader IT and Business Process Outsourcing (BPO) sectors be impacted? Will the labor impacts of AI remain external to banks and other implementing organizations, or will the internal implications grow with time as companies increasingly adopt and integrate AI? If the impacts are externalized, who will take responsibility for the

people who are affected by those externalities? Other questions pertain more broadly to the societal outcomes that AI systems might facilitate: Will AI increase financial access by scaling banking services to mobile users, or might it deepen the accessibility gap due to implementation challenges with language translation and unequal access to internet banking? Should consumers be notified when they are interacting conversationally with an automated system? Further research involving the perspectives of both workers and consumers could be particularly valuable as we collectively come to understand the nature of AI-use in the real world.

Appendix

Definitions and terms used

While we acknowledge that there is no consensus on the definition of terms such as AI and automation, we would like to explain how these terms are used in the compendium:

Artificial intelligence/AI is a notoriously nebulous term. Following the [Stanford 100 Year Study on Artificial Intelligence](#), we embrace a broad and evolving definition of AI. As Nils J. Nilsson has articulated, artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment. (Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge, UK: Cambridge University Press, 2010).

Our definition of **automation** is based on the classic human factors engineering definition put forward by Parasuraman, Sheridan, and Wickens in 2000: <https://ieeexplore.ieee.org/document/844354>, in which automation refers to the full or partial replacement of a function previously carried out by a human

operator.⁷ Following Parasuraman et al.'s definition, levels of automation also exist on a spectrum, ranging from simple automation requiring manual input to a high level of automation requiring little to no human intervention in the context of a defined activity.

Explainable AI or **Explainability** is an emerging area of interest in communities ranging from DARPA to criminal justice advocates. Broadly, the terms refer to a system that has not been “black-boxed,” but rather produces outputs that are interpretable, legible, transparent, or otherwise explainable to some set of stakeholders.

In this compendium, a **model** refers to a simplified representation of formalized relations between economic, engineering, manufacturing, social, or other types of situations and natural phenomena, simulated with the help of a computer system.

⁷ Our definition draws on the classic articulation of automation described by Parasuraman, Sheridan, and Wickens (2000): <https://ieeexplore.ieee.org/document/844354>



PARTNERSHIP ON AI