Journal of Ethics and Emerging Technologies

IEET

Practice Paper

Automation and Augmentation: Practical Ethics of AI for Productivity in Retail Banking

Albi Nani [1]*

[1] Independent; anani.academia@gmail.com
* Correspondence: anani.academia@gmail.com

Abstract: AI is increasingly used as a tool by firms to enhance productivity, but its success has been accompanied with ethical dilemmas concerning the employment of people involved in the tasks touched by AI. The popular narrative states that AI will replace human workers, substantiated by layoffs where AI-driven efficiencies have been cited as a primary catalyst. In practice, the relationship between AI and labor ethics is more nuanced. This practice paper makes two contributions by commenting on industry practices concerning AI ethics. First, using the example of retail banking, a model is provided of how firms categorize AI use-cases based on time and cognitive effort, resulting in either automation or augmentation. Second, the ethical dilemmas of both options are explored; automation does not drive workforce reduction as much as popular speculation would suggest, while the real ethical risk of augmentation is related to task circumvention and knowledge loss.

Keywords: labor, AI, AI ethics, automation, banking, productivity, AI productivity, artificial intelligence, augmentation

Citation: Nani, Albi,. 2025. Automation and Augmentation: Practical Ethics of AI for Productivity in Retail Banking. *Journal of Ethics and Emerging Technologies* 36: 1. https://doi.org/10.55613/jeet.v36i1.200

Received: 22/10/2025
Accepted: 18/03/2026
Published: 24/03/2026

Publisher's Note: IEET stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction: AI, Ethics, and Productivity

As AI is more frequently used in large organizations to generate business value, the use of AI systems to increase productivity has become particularly popular (Bainbridge, 2025). In a business context, we define productivity as 'the ratio of output versus input' (Kobiella et al., 2024), though inputs and outputs can vary depending on the specific industry or work involved. Generally speaking, AI has been successfully leveraged for productivity gains; research by the OECD claims that AI can positively impact productivity anywhere from 10-56%, depending on the specific task measured (Filippucci et al., 2024).

However, the increase in productivity has often been accompanied by ethical concerns for labor, either drastically transforming the nature of work (Bankins & Formosa, 2023), or potentially replacing humans in the workforce entirely (Ramachandran et al., 2024). While these concerns are valid, they are often conceptualized theoretically in academic circles without opportunities to be examined in practice. Therefore, this practice paper analyses these concerns in practice and asks the following research question: *how do organizations categorize the impact of AI on specific tasks, and what are the ethical concerns associated with each category?* To provide more specificity, the examples provided in this paper will be specific to the retail banking industry in Canada.[1]

[1] Retail banking was selected as the case for analysis due to the author's professional experience as an AI strategist and consultant at a chartered financial institution in Canada; ensuring that the insights provided reflect business reality.

First, this paper proposes a descriptive model for AI deployment, providing insight into how organizations determine whether a task should be augmented or automated by AI. This model also describes how AI produces productivity gains when deployed against certain tasks. Next, the ethical implications of augmentation and automation are discussed, providing perspective on how organizations navigate these concerns in practice, before concluding with a summary of the overall argument.

2. The AI Deployment Model

2.1. Productivity in Retail Banking

Retail banking refers to the collection of banking products and services that are offered to most individuals and small businesses.[2] Retail banking services are offered through a variety of channels, including digital (i.e. mobile applications and online banking), and traditional banking center branch locations. Despite the convenience of digital options, consumers still have an affinity for visiting banking centers; a KPMG survey revealed that in Canada, 34% of clients visited a branch at least once a month, and a further 41% visited at least once a year (2024). Therefore, optimizing productivity in banking centers is a strategic priority for retail banks.

The staff within branches are typically a mix of service and sales workers;[3] the former help clients transact (e.g. pay bills, withdraw funds, etc.), while the latter encourage clients to open new accounts or deposit more funds in the bank. Generally speaking, the most valuable expenditure of time is meeting with clients, allowing staff to understand, address, and fulfill their needs, consequently strengthening the bank's relationship with a client. However, face time with clients is often limited by the time spent on administrative and manual work such as conducting client research, filling out forms, and taking notes, among other activities (McKinsey, 2024). Optimizing productivity in a banking center means maximizing client-facing time, while minimizing the time and resources spent on administrative work. AI provides opportunities to bolster productivity in retail banking, but the exact tactics and solutions used will vary depending on the type of tasks being impacted.

2.2. Assessing Tasks

While it is generally agreed that AI can improve productivity across a variety of tasks (Czarnitzki et al., 2023), not all tasks are impacted in the same manner. A simple method to assess the impact of AI on a task is to characterize them based on their volume and type of input. Some tasks may require more time to complete (e.g. data entry), while others necessitate more intellectual complexity (e.g. offering financial advice). To understand the needs of a specific task, they can be assessed against these inputs. Time commitment describes the amount of time required to complete a task, while cognitive complexity describes how much nuance and thoughtfulness is required for task execution. It should be noted that these qualities are not mutually exclusive; a task can require both a substantial time commitment and cognitive effort to complete.

Measuring these variables accurately ensures that the task is correctly characterized. As a more tangible variable, time commitment is easier to measure. In many cases this is achieved through a time study of various tasks following traditional Taylorist practice (Börnfelt, 2023). Cognitive effort is more abstract, and therefore more difficult to measure.

[2] This is in opposition to private wealth management, which offers customized solutions to high net-worth individuals, and commercial banking, which addresses the needs of large businesses.

[3] Service staff may hold titles such as Customer Service Representatives or Tellers, while sales staff often hold the title of Financial Advisor.

Moreover, the immense variety of tasks means that a standardized assessment method is typically not a feasible option. Oftentimes, firms will resort to ranking the complexity of tasks relative to one another using worker feedback, providing an approximation of which tasks have the highest cognitive demands.

2.3. Deploying AI

Once tasks have been assessed for their requisite inputs, they can be plotted against two axes as illustrated in **Figure 1**. The plot is divided into three areas that describe how AI transforms a given task: automation, augmentation, and orchestration. Though similar models have been used in industry, no such model has been introduced yet in academic discourse.[4] The presented model is a propositional framework, synthesizing industry practices that would require further empirical evidence to validate.

Though cognitive load and time commitment are separate axes, they are inherently related to one another; a complicated task will naturally take more time to complete than a simple one. To reconcile these differences, the model is best understood in relative terms rather than absolute measurements. For instance, if two tasks took the same amount of time, then differences in cognitive load could determine if automation or augmentation is the better solution.

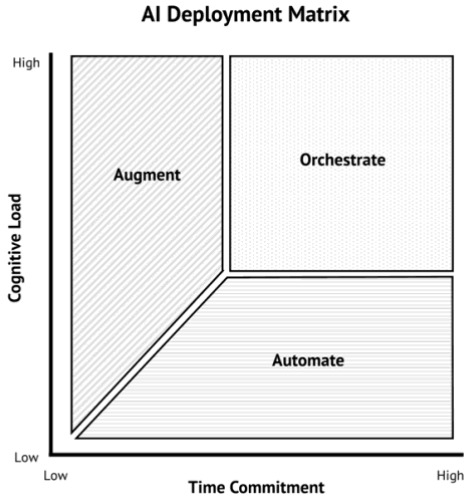

Figure 1: AI Deployment Model[5]

When a task requires low cognitive effort and high time commitment, then AI is best used for automation. In lay terms, this is work considered to be easy but time-consuming when done by a human. In retail banking, this includes tasks such as entering client data into systems, creating post-meeting notes, and writing e-mails to clients. These tasks are ripe for automation due to their simplicity; with less complexity to consider, firms can reap the benefits of a machine completing the work faster while minimizing errors or inaccuracies that would otherwise hamper productivity (Simkute et al., 2025). The Canadian Imperial Bank of Commerce (CIBC) did exactly this (2025), giving employees access to an

[4] Similar models have been proposed by corporate thought leaders on LinkedIn (https://www.linkedin.com/pulse/automate-augment-approve-assist-troy-latter-xhqjc/), or in privately accessible workshops held by technology firms such as Amazon. The specific iteration proposed in this paper is unique but draws from these consulting-type resources.

[5] If a task is rates as low in time commitment and cognitive effort, it can be either augmented or automated based on the specific work involved. For this reason, the plot is not a true matrix as there is no distinctive solution for tasks with low time commitment and cognitive effort.

internally-developed Gen-AI application that saw significant use in drafting e-mails, presentations, and other communications. Productivity is achieved through automation by lowering the time input into task completion, allowing that time to be redistributed to client-facing activity.

If a task requires low time commitment but has relatively high cognitive complexity, then AI is most effective at augmenting a human user's completion of that task. AI enhances the user's knowledge and capabilities with facilitated access to information and complex insights. Examples within retail banking include supporting financial advisors with financial advice and product recommendations for clients, and summarizing data as part of client research. Toronto Dominion (TD) Bank (2025) implemented a Gen-AI knowledge search application that emulated this framework, surfacing information to contact center agents that provided answers to client inquiries. By eliminating manual search and letting agents focus on serving clients, TD was able to reduce call hold times by 15%. In contrast to automation, augmentation impacts productivity by improving the output- often qualitatively- of the services offered (Shaunn Mattingly et al., 2025). The cognitive and creative demands of these tasks often necessitate that a human still be in control, but AI can expedite access to information that can enrich the user's work.

Some tasks require a combination of both high time commitment and cognitive load to complete, meaning AI can both augment the user and automate the task to successfully generate productivity gains. This combination is termed *orchestration*, as the AI is effectively guiding the user through task completion to decrease the time input while simultaneously improving the quality of the output. Most administrative work is still automated, but certain elements of the role are augmented to maintain a degree of human involvement. In retail banking, orchestration is used to guide workers through certain *journeys*- a series of actions and tasks that result in a desired outcome. For example, a bank would orchestrate a journey for a financial advisor to use when acquiring and onboarding a new client; the sales conversation may be augmented with product recommendations for a client, but the backend account setup would be automated.

While one could argue that orchestration is simply a combination of automated and augmented tasks, there needs to be consideration for how these tasks fit together in a broader workflow, hence the importance of the journey view. As AI becomes increasingly sophisticated and capable of completing an array of tasks, it will be possible to collect enough of these AI-enhanced tasks to redefine entire processes. For instance, if a bank deployed an AI system to provide investment recommendations to clients, and another AI agent that automated the fulfillment of investment products, the combination of these two orchestrates the process of selling investments to the client. This would not be the case if the two tasks were unrelated to one another.

In summation, orchestration deploys AI in whatever capacity necessary to optimize productivity across a journey of related tasks. However, questions remain as to how the people originally performing these tasks are impacted by AI-driven productivity.

3. Ethical Implications for AI and Productivity

While business will naturally prioritize operational benefits from implementing AI use-cases, these decisions inevitably impact the people originally performing these tasks. Therefore, there are ethical questions present related to the employment and role of workers in organizations implementing AI, though the exact impacts differ between automation and augmentation.

3.1. Ethical Implications for Automation

Concerns of automation leaving individuals out of work date back to older industrial revolutions, where machines capable of mass production were used to improve efficiency on factory floors (Noponen et al., 2024). AI, and particularly Gen-AI, has expanded the scope of tasks that can be completed autonomously by a machine; for instance, a sufficiently advanced chatbot is capable of solving a client's servicing request (Shaunn Mattingly et al., 2025)- in the banking context this might be a credit card limit increase.

Critics of AI-driven automation outline how AI threatens to eliminate the jobs of the humans that would have previously completed the automated work (Tilleczek, 2025). To a degree, these concerns are valid, as many firms have taken the opportunity to downsize their workforce, citing AI as a specific catalyst for enabling workforce reduction (Scheiber, 2025). However, these critiques often miss some of the nuance involved with AI-related workforce reduction.

One such nuance is that mass lay-offs are rarely the sole result of AI. It would be extremely risky for firms to terminate a significant portion of their workforce at once without giving their operations time to adjust; for a retail bank, these risks include leaving clients underserved, ultimately losing their business and causing reputational damage. For organizations as risk-averse as Canadian banks,[6] layoffs would not be driven by AI.

Instead, it is more common for AI-driven workforce reduction to occur over time as a result of regular employee attrition. This works especially well for contact center or similar service roles, which have a high degree of natural attrition. Essentially, the organization distributes AI capabilities to employees and experiences an overall increase in productivity; as workers naturally attrit, the organization may not replace the role due to the newfound productivity gains. This approach mitigates the risk of operational disruption while still letting an organization experience the productivity benefits of AI. This aligns with the approach taken by Scotiabank when the organization laid-off 3,000 workers in 2025, citing opportunities for rationalization (Karim, 2026). However, AI was not mentioned as an explicit reason for layoffs, which contrasts lay-off announcement from other firms, namely in Big Tech, made in 2025.

While it would be disingenuous to ignore that other industries and firms are experiencing mass layoffs in which AI-driven productivity is the motivating factor, it is unwise and incorrect to accept this as a universal truth across all businesses. Ethical criticisms of big tech are not ubiquitous to other industries. Many retail banks globally are infamous for outdated technical infrastructure (Hayretci & Aydemir, 2021), and do not have the capability to instantly implement AI and transform their workforce overnight. When considering the ethical implications for AI automation on labor, it is critical to look at the specific organizational context in which these transformations are occurring.

3.2. Ethical Implications for Augmentation

Where automation reduces the input time required to complete work, augmentation changes the nature of the work itself- from a human perspective- to enhance the quality of the output. This often manifests into arguments that AI makes work more meaningful for humans by leaving them the tasks that provide the most engagement and value (Bankins & Formosa, 2023). In the context of banking, this would mean affording a financial advisor more time to meet with clients by taking away administrative work.

[6] Banks are considered a highly-regulated industry, meaning that they are naturally more risk averse as they must comply with more regulations.

While the exact requirements for meaningful work are subjective, how workers feel about the work they have accomplished may be a suitable proxy indicator. To that end, work by Kobiella et al. (2024) demonstrates that workers often feel more accomplished when using AI tools, citing "greater creative output" and "efficient completion of tasks". Accomplishment may also stem from better aligning work time with compensation incentives. Financial advisor compensation can often depend on their funds managed relative to their peers;[7] if AI can help by augmenting their sales tactics, the work can be described as more meaningful because it substantially helps advisors increase their likelihood of more compensation.

Rather, the ethical dilemma related to augmentation stems from it being too useful; at a certain point, augmentation may turn into task circumvention. For example, if AI proves to be a sufficiently effective tool for client research, then an advisor may rely on the AI tool to do the entirety of the research instead of using it as a supporting feature. Over time, this means that an advisor is relying less on their own knowledge and memory and instead turning to AI tools whenever they need information on their own clients. Essential skills are circumvented entirely as they are gradually offloaded to a AI tools (Tilleczek, 2025). This is especially concerning should an AI system be sufficiently trained to complete tasks that require certification for humans to complete.

Financial advisors must obtain their IIROC or MFDA licenses before they are allowed to hold their position and advise clients. Obtaining these licenses involves completing courses that teach candidates specialized knowledge about financial products, risk, and investments. If in the future, advisors can simply leverage an AI to conduct financial planning for clients, then they risk losing the knowledge required that made them certified advisors initially. This argument can be applied across most roles with specialized knowledge; overreliance on AI for augmentation eventually becomes a loss of knowledge as tasks are outsourced to a machine. Admittedly, this risk is more of a future concern than a current problem; AI applications are not yet mature or sophisticated enough in many banking ecosystems to autonomously complete complex financial plans and investment strategies. However, given the rapid rate of AI development, it would be wise to plan ahead and proactively address these risks. The Royal Bank of Canada (RBC) is already deploying AI assistants for capital markets research (2025); deploying similar technologies at scale for retail banking is likely feasible in the near future.

Instead of AI augmenting human labor, over time the opposite will be true as AI systems become more sophisticated through human-computer interactions. While it is popular to implement a human-in-the-loop to oversee outputs as a risk mitigation strategy (Natarajan et al., 2025), what good is a human if they lack the knowledge to police AI systems after using these same AI tools to circumvent their work? Organizations should recognize AI-induced knowledge loss as a risk and find ways to maintain the skill and education of their workforce, even when AI incentivizes the disposal of that knowledge.

4. Conclusion

The use of AI to create productivity gains is effective but will also have profound ethical impacts on the way that work occurs. This practice paper contributes to developing our understanding of AI and ethics by (1) providing a framework for how organizations think about AI and productivity, and (2) exploring the nuances of how differently AI impacts certain tasks and the roles associated with them. While ethical concerns of workforce reduction and deskilling labor are valid, AI is not the catalyst of these outcomes. Rather, AI accelerates the fulfillment of these results by substantially increasing productivity in

[7] Funds managed refers to the total investible assets (e.g. cash, stocks, GICs. etc.) held by a client or a portfolio of clients.

organizations. AI is a potent productivity tool, but ultimately it is the context of its deployment that should be the subject of ethical analyses.

Moreover, this paper also provides two practical contributions to help practitioners navigate ethical decisions involving AI and productivity. First (1), while AI can contribute to more productive business models, it is rarely capable of replacing a workforce on its own. Productivity gains should be experienced for sustained periods of time before drastic decisions are made about a firm's workforce. Second (2), the outsourcing of knowledge to AI systems presents a real risk for organizations whose offerings are reliant on human expertise. These organizations should develop strategies for maintaining this knowledge in their workforce, preventing overreliance on AI tools.

Funding: This research received no external funding.

Acknowledgments: I would like to thank the editors and reviewers for their support and commentary, as their contributions were critical in developing this work.

Conflicts of Interest: The author declares no conflict of interest.

References

Bainbridge, W. (2025). The Future of Worker Enhancement. *Journal of Ethics and Emerging Technologies*, *35*(1), Article 1. (2016+ US). https://doi.org/10.55613/jeet.v35i1.150

Bankins, S., & Formosa, P. (2023). The Ethical Implications of Artificial Intelligence (AI) For Meaningful Work. *Journal of Business Ethics*, *185*(4), 725–740. https://doi.org/10.1007/s10551-023-05339-7

Börnfelt, P.-O. (2023). Taylorism and Fordism. In P.-O. Börnfelt (Ed.), *Work Organisation in Practice: From Taylorism to Sustainable Work Organisations* (pp. 25–37). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-21667-1_2

CIBC. (2025). *Transformation at scale: CIBC's strategic leap into Enterprise Generative AI*. https://www.cibc.com/en/about-cibc/future-banking/ai/enterprise-generative-ai.html

Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, *211*, 188–205. https://doi.org/10.1016/j.jebo.2023.05.008

Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). The impact of Artificial Intelligence on productivity, distribution and growth. *OECD Artificial Intelligence Papers*, *15*. https://doi.org/https://doi.org/10.1787/8d900037-en.

Hayretci, H. E., & Aydemir, F. B. (2021). A Multi Case Study on Legacy System Migration in the Banking Industry. In M. La Rosa, S. Sadiq, & E. Teniente (Eds.), *Advanced Information Systems Engineering* (pp. 536–550). Springer International Publishing. https://doi.org/10.1007/978-3-030-79382-1_32

Karim, N. (2026, January 8). Scotiabank let 3,000 people go late last year, CEO says. *Financial Post*. https://financialpost.com/fp-finance/banking/scotiabank-laid-off-3000-people-2025

Kobiella, C., Flores López, Y. S., Waltenberger, F., Draxler, F., & Schmidt, A. (2024). "If the Machine Is As Good As Me, Then What Use Am I?" – How the Use of ChatGPT Changes Young Professionals' Perception of Productivity and Accomplishment. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24*, 1–16. https://doi.org/10.1145/3613904.3641964

KPMG. (2024). *Canadian banks must innovate their branch networks, KPMG survey shows—KPMG Canada*. https://kpmg.com/ca/en/home/media/press-releases/2024/06/canadian-bank-branches-need-innovation-kpmg-poll-shows.html

McKinsey. (2024, December 9). Extracting value from AI in banking: Rewiring the enterprise. *McKinsey Financial Services*. https://mckinsey.com/industries/financial-services/our-insights/extracting-value-from-ai-in-banking-rewiring-the-enterprise

Munn, L. (2023). The uselessness of AI ethics. *AI and Ethics*, *3*(3), 869–877. https://doi.org/10.1007/s43681-022-00209-w

Natarajan, S., Mathur, S., Sidheekh, S., Stammer, W., & Kersting, K. (2025). Human-in-the-loop or AI-in-the-loop? Automate or Collaborate? *Proceedings of the AAAI Conference on Artificial Intelligence*, *39*(27), 28594–28600. https://doi.org/10.1609/aaai.v39i27.35083

Noponen, N., Feshchenko, P., Auvinen, T., Luoma-aho, V., & Abrahamsson, P. (2024). Taylorism on steroids or enabling autonomy? A systematic review of algorithmic management. *Management Review Quarterly*, *74*(3), 1695–1721. https://doi.org/10.1007/s11301-023-00345-5

Ramachandran, Kk., Raju, V., Karthick, K., Lakshmi, Baba Gnanakumar, P., & Deepa, M. (2024). Rise of AI: Prediction of Job Replacements Based on the Evolution of Artificial Intelligence and Robots Intensification. *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)*, 1–6. https://doi.org/10.1109/ACCAI61061.2024.10602094

RBC. (2025, June 9). *How RBC is using AI to transform capital markets research | RBCCM*. https://www.rbccm.com/en/insights/2025/06/how-rbc-is-using-ai-to-transform-capital-markets-research/

Scheiber, N. (2025, July 7). Which Workers Will A.I. Hurt Most: The Young or the Experienced? *The New York Times*. https://www.nytimes.com/2025/07/07/business/ai-job-cuts.html

Shaunn Mattingly, E., Kroes, J. R., & Manikas, A. S. (2025). AI as an emerging trend for managing employee efficiency in the retail and services industries. *Business Horizons*. https://doi.org/10.1016/j.bushor.2025.06.005

Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2025). Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction. *International Journal of Human–Computer Interaction*, *41*(5), 2898–2919. https://doi.org/10.1080/10447318.2024.2405782

TD Bank. (2025). *TD customer service in the age of AI*. https://stories.td.com/ca/en/article/td-customer-support-agent-ai-tool

Tilleczek, W. (2025). Practice or Product? Labour, Training, and the "Ethics" of AI. *Global Labour Journal*, *16*(2). https://doi.org/10.15173/glj.v16i2.5861