PROJECT 'EMPIRE'

Final Capstone Paper

Challenges in calculating the ROI to Generative AI for financial institutions

Abstract

This paper presents the merit in creating a framework to determine the Return on Investment (ROI) of Generative AI. As a formula for calculating the ROI on regular investments is rather simple, calculating the ROI for Generative AI should not have been too challenging. Structurally, we considered a private-benefit approach followed by a public-benefit approach. The former analyzed possibilities through which a formula that could demonstrate improvements in Net Income or Profit by employing Generative AI could be designed. Specifically, we considered an approach that measures potential improvement through COGS or SG&A, a historical approach that uses data from firms that have already implemented Generative AI, and a decision tree model. We have concluded that calculating a framework to determine the ROI on Generative AI is not possible at the moment as not enough firms have implemented the technology, the technology is too nascent, and the associated data is too scarce. Any actor attempting to create an ROI for Generative AI at this time will face the same challenges outlined in this paper and reach the same conclusion. Future research could be focused on individual business functions and specific applications to create a framework solely for a specific application. The culprit of this approach is that it would only apply to a specific application within a specific firm.

Table of Contents

Introduction

Calculating the Return on Investment (ROI) is imperative to valuing any capital expenditure. Recent developments in Artificial Intelligence (AI) have clearly illustrated its capabilities in increasing productivity per employee and cutting costs. This feature has made AI attractive amongst firms seeking to increase their profits by employing this technology. Thus, the research we conducted had the following goal in mind; developing a framework to determine the ROI of Generative AI.

A common theme throughout this paper is the fact that creating a framework to determine the ROI on Generative AI should theoretically be straightforward. ROI is calculated using a well-known, simple formula, and should not be too challenging to translate into a formula for measuring the ROI of Generative AI. Nevertheless, what seemed to be straightforward theoretically proved to be challenging and infeasible, practically.

This paper presents and evaluates our research chronologically, initially setting out to find the solution, not the challenges. It will naturally follow a narrative illustrating the challenges with each approach we pursued. Thus we will intrinsically conclude that the current data availability renders it impossible, for everyone who will try, to develop such a framework at this time.

To obtain a holistic and critical perspective on the matter we will utilize a comparative methodology that will evaluate the potential private benefit in Part 1 and the potential social benefit in Part 2. In Part 1, we will showcase how ROI is typically calculated, and juxtapose this to how ROI should be calculated for Generative AI. Subsequently, we will present a literature review which will allow us to narrow down an industry and firms. The core of this section explains what could possibly act as a numerator for our function. Finally, we will present the challenges and an alternative approach. As the ROI for the private sector is measured in terms of Net Income and Profit, the bulk of this paper will regard the first section. The second section will regard the Social Benefit/the prospect of Generative AI. We will discuss the idea of a social utility function and evaluate the advantages and disadvantages of Generative AI from a social perspective as the technology will likely become a public good and will be difficult to generate revenue once so.

Prior to commencing with the analysis, it is imperative to define AI and Generative AI. AI is essentially a behavior used by computers to simulate certain productive life activities that only people can perform. This includes the ability to learn, reason, perceive, understand, plan, make decisions, etc. The goal of AI is to enable computer systems to mimic human thought processes in order to solve a wide range of complex problems, thereby enabling autonomous action and intelligent decision-making. Generative AI is a type of artificial intelligence whose main function is to automatically generate corresponding content, such as text, images, videos, etc., based on inputs, not just simply analyzing existing data or performing specific tasks. Generative AI systems are usually based on machine learning and deep learning techniques, and are trained on large-scale data to obtain models that can be used in different fields, such as technology, finance, education, etc.

PART I

Calculating ROI

ROI can be regarded as a financial metric widely used to measure the probability of gaining a return from an investment. It is a ratio of the profit or loss made by an investment relative to its initial cost. The overarching principle here is that the benefits derived from a particular investment needs to surpass the company's hurdle rate, thereby ensuring the investment is justifiable.

In simple terms, the formula for ROI is Net Return divided by Cost of Investment. The numerator, Net Return, represents the total profit received from an investment. This could be identified through metrics such as an increase in revenue or a reduction in cost. The denominator, Cost of Investment, represents the total amount spent.

$$Return\ on\ Investment\ (ROI) = \frac{Net\ return}{Cost\ of\ Investment} \times 100$$

ROI is a prevalent financial metric employed extensively in determining the feasibility of proceeding with a particular investment. Consequently, we are examining this particular metric for potential investments in the domain of Generative AI.

Calculating ROI from Generative AI

As Generative AI is a term that represents a myriad of technologies, processes, and use cases, determining an ROI is complex and not a one-size-fits-all for every company. Thus, the task lies in identifying the crucial metrics that can measure the net return and cost of investment.

Calculating this ROI presents unique challenges, particularly when considering risk-adjusted returns. Risk-adjusted return will focus on various potential costs associated with Generative AI. These include the potential for hallucination, opportunities for cyberattacks, data breaches, improperly trained models, and copyright concerns.[1] Incorporating these risks into ROI calculations involves a deep understanding of the legal, financial, and ethical implications of the generated contents.[2]

The NIST AI Risk Management Framework (RMF) offers valuable guidance in this area, emphasizing the importance of risk measurement across different stages of the AI lifecycle and in real-world settings, as well as the significance of AI trustworthiness in terms of validity, safety, security, resilience, accountability, transparency, explainability, interpretability, privacy, and fairness. These factors complicate the ROI calculation for Generative AI, adding layers of complexity beyond traditional investment analysis.[3]

[1] OliverWyman, “How Generative AI is transforming business and society: The good, the bad, and everything in between”, 2024.
[2] Harvard Business Review, “Managing the Risks of Generative AI”, 2023.
[3] NIST, “Artificial Intelligence Risk Management Framework”, 2023.

Literature review

As an organization matures in its adoption and utilization of AI technologies, the potential benefits to be derived from AI deployment may increase significantly. On average, companies realize an ROI of 5.9% from AI projects, with leading organizations achieving an ROI of 13% on AI projects. Key factors contributing to the development of AI capabilities include vision and strategy, AI operating models, AI engineering and operations, data and technology, talent and skills, and culture and adoption.[4] These numbers have been derived from general AI projects and are not specifically targeted at Generative AI.

Generative AI, when utilized within the bounds of its capabilities, may enhance the productivity of high-skill workers by approximately 40%. However, exceeding the boundaries of Generative AI capabilities may lead to a decrease in productivity by an average of 19%. These findings underscore the importance of understanding the limitations of such technologies for optimizing workers' productivity.[5]

The benefits stemming from AI implementation manifest in both tangible and intangible forms. Hard benefits include measurable outcomes such as time savings, productivity enhancements, cost reductions, and revenue growth, whereas soft returns encompass improved user experience, skill retention, and organizational agility.[6]

The potential impact of Generative AI is substantial, with estimates suggesting it could contribute $4.4 trillion annually to the global economy, primarily through significant productivity gains across various industries. Amongst all industries, retail and consumer-packaged goods, banking, pharmaceuticals, and medical products are anticipated to derive significant benefits from their adoption.[7]

Generative AI presents opportunities across diverse business functions, with approximately 75% of its potential value stemming from customer operations, marketing and sales, software engineering, and research and development (R&D) when 63 use cases of Generative AI across 16 business functions have been studied.[8]

By automating tasks that currently occupy 60-70% of employee time, Generative AI is expected to enhance customer interactions, craft content for marketing and sales, and generate executable codes based on natural language prompts.[9] However, realizing these benefits necessitates investments in skill development and support for workers transitioning between job activities or roles.

Projections indicate that Generative AI could boost annual labor productivity growth rates by 0.1 to 0.6 percent until 2014, contingent upon the pace of technology adoption and effective reallocation of workers' time to alternative tasks. Combined with other automation technologies,

[4] IBM, "Generating ROI with AI", 2023.
[5] MIT, "How Generative AI Can Boost Highly Skilled Workers' Productivity", 2023.
[6] PWC, "Solving AI's ROI Problem. It's Not That Easy", 2021.
[7] McKinsey & Company, "The Economic Potential of Generative AI: The Next Productivity Frontier", 2023.
[8] McKinsey & Company, "The Economic Potential of Generative AI: The Next Productivity Frontier", 2023.
[9] McKinsey & Company, "The state of AI in 2023: Generative AI's breakout year", 2023.

the cumulative impact could produce productivity growth rates of 0.2 to 3.3 percent annually. Nevertheless, facilitating economic growth through enhanced labor productivity requires investments in skills training and supporting occupational transitions for workers.[10]

While offering potential opportunities, Generative AI also presents legal, financial, and ethical risks that may pose threats due to the inadequacy of regulatory frameworks and early-stage development.[11] Potential dangers may predominantly include AI hallucinations, data breaches, cyberattacks, improperly trained models, and copyright issues. Managing these risks requires robust governance frameworks and ethical AI practices to mitigate possible threats.[12]

Given the varying degrees of integration across industries and the nascent stage of Generative AI development, strategies for AI deployment need to be tailored within organizations to reap the greatest benefits.

- **ROI:** Best-in-class reap a 13% ROI on AI projects, more than twice the average of ROI of 5.9% (IBM)
- **Labor productivity:** Improvement in a worker's performance (high-skill) by as much as 40% compared with workers who don't use it (MIT Sloan, 2023)
- **Automation:** Automate tasks that absorb 60-70% of employee's time (Mckinsey, 2023)
- **Firm values:** Firms with higher exposure of their labor force to Gen AI outperformed firms with lower exposures by over 40 basis points in daily excess returns on portfolios (NBER, 2023)
- **Hard and soft returns:** Hard returns include time savings, productivity increase, cost savings, and revenue increase, while soft returns include better experience, skill retention, and organizational agility (PwC, 2021)
- **Efficiency in business functions:** Majority of the value that is brought by AI falls under the following four areas: Customer operations, marketing and sales, software engineering, and R&D (Mckinsey, 2023)
- **Gen AI adoption:** More than 70% of companies are experimenting with Gen AI, but less than 20% is willing to spend more on it (Deloitte, 2023)
- **Risks:** Considering the legal, financial and ethical implications of the generated content is integral in the usage of AI, for cases where the results are inaccurate (Harvard Business Review, 2023)
- **Sustainability:** LLMs have hundreds of of billions of parameters and use a lot of energy and water to train them (Harvard Business Review, 2023)

Chart 1: Key metrics associated with the literature review

Industry selection

Generative AI is poised to revolutionize various industries, with certain sectors expected to experience a more pronounced impact. Our literature review suggests that the impact of Generative AI will vary depending on the depth of integration with business functions. This variability underscores the necessity for tailored AI strategies based on specific industry requirements.

In this nascent phase of technology adoption, and with varying degrees of penetration across industries, data availability for calculating the ROI remains limited. Moreover, Generative AI presents vast opportunities across diverse business functions. Thus, challenges in ROI measurement arise due to differing exposure levels and metrics. Hence, identifying certain

[10] McKinsey & Company, "The Economic Potential of Generative AI: The Next Productivity Frontier", 2023.
[11] Harvard Business Review, "Managing the Risks of Generative AI", 2023.
[12] Oliver Wyman, "How Generative AI Is Transforming Business and Society", 2024.

criteria is crucial to focus the analysis on industries that are most likely to benefit from Generative AI.

The selection of industries is guided by three key criteria: market capitalization, potential revenue from AI adoption, and integration capability. Market capitalization, primarily focusing on S&P 500 companies classified under the Global Industry Classification Standard, served as the first criterion. This metric reflects a company's total stock value, providing insights into its market size and influence. By targeting S&P 500 companies, the analysis hones in on industry leaders that spearhead technological advancements, including Generative AI implementations.

The second criterion, potential revenue from AI adoption, directly quantifies the financial benefits a company might achieve through Generative AI. The literature suggests that across different use cases, Generative AI has the potential to generate up to $4.4 trillion in value[13] across industries. Its precise impact will depend on a variety of factors, such as the mix and importance of different functions, as well as the scale of an industry's revenue. Research highlights several sectors such as banking, high tech, and industrials (Chart 1) as significant beneficiaries. These sectors could witness substantial impacts in terms of revenue generation due to the integration of Generative AI technologies. With its current capabilities, Generative AI holds the potential to deliver substantial value within these sectors.

Finally, integration capability serves as the third indicator, emphasizing the seamless integration of AI solutions into existing organizational systems and processes within an organization. A smooth integration ensures that AI tools complement and enhance the current operations rather than causing disruptions. Measuring integration capability could involve assessing compatibility with existing technology, ease of implementation, and the level of change management required. As reported by Bain & Company[14], companies with high levels of automatable and augmentable roles see more significant benefits from Generative AI. This is because these organizations are inherently more adaptable to technological integration, making them better positioned to capitalize on AI advancements. Therefore, a metric such as the percentage of employees in roles with high automation or augmentation potential, computed by Bain&Co in the above-mentioned report is taken as a proxy to measure the potential scale of AI integration capabilities.

[13]McKinsey & Company, "The Economic Potential of Generative AI: The Next Productivity Frontier", 2023.
[14]Bain & Company, "AI Investors: Act Fast, Act Wisely", September 2023

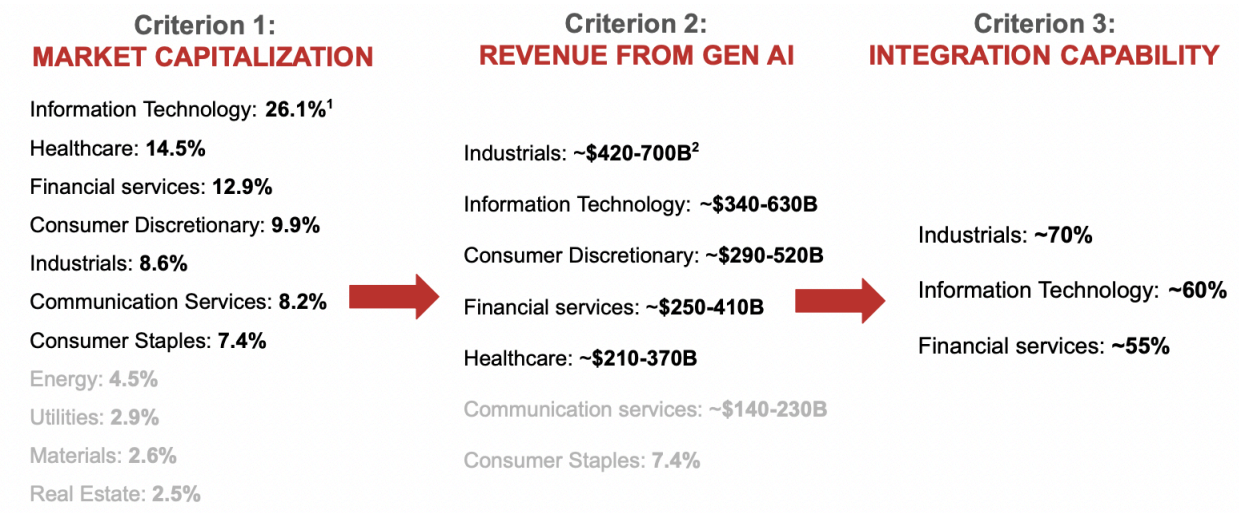

Chart 2: Industry selection

In conclusion, based on market capitalization, potential revenue from AI adoption, and integration capability criteria, Generative AI offers the most opportunities for the financial industry and is set to redefine it. By enhancing customer service, automating routine tasks, and bolstering fraud detection, this technology is projected to generate an additional annual revenue of $200 billion to $400 billion, promising significant economic benefits and operational efficiencies for the sector.

Generative AI Use Cases in the Financial Industry

As we delved deeper to understand the uses of Generative AI in the Financial Industry, we sought first-hand information from industry practitioners who shared insights in informal discussions, albeit within their professional constraints. The Financial Industry has been increasingly integrating Generative AI into its operations, particularly in internal services and team communications. Engineers are leveraging Generative AI to streamline processes, such as answering questions related to coding and addressing technical issues, enhancing efficiency, and reducing response times[15]. One of the primary reasons for this adoption is the perceived low risk associated with utilizing Generative AI in these internal functions. Additionally, there is a focus on training the AI to handle specific information and ensuring quality assurance, with the potential for deployment in client-facing solutions. Despite the potential benefits, as of now, there is hesitance in deploying Generative AI for client interactions due to concerns about trust. Companies are wary of the perceived high risk, as clients prefer human interactions, citing a trust factor about Generative AI that has not yet been fully established.

Moreover, the industry sees the potential for training Generative AI models with specific information, such as Master Service Agreement (MSA) terms and related common questions, to ensure quality assurance in client-facing solutions. This approach holds promise for enhancing client interactions by providing accurate and standardized responses to inquiries, ultimately

[15] Primary interview with employees at multinational investment bank and financial services company, April 2024.

improving the overall client service experience. However, companies have approached the deployment of Generative AI in client-facing roles cautiously, citing concerns over perceived high risk and the enduring preference of clients to interact with human agents, underscoring the pivotal role of trust in financial transactions.

Here, we look at the findings from advisory services providers like JP Morgan Chase and Morgan Stanley, who have conducted studies revealing significant productivity gains through the use of Generative AI in advisory services. These AI-empowered advisory assistants have demonstrated a 21% increase in productivity compared to traditional methods[16]. Such findings underscore the growing recognition within the financial industry of the transformative potential that Generative AI promises in enhancing advisory services. Despite initial reservations, these successes are likely to encourage further exploration and adoption of Generative AI across a broader spectrum of financial services, potentially reshaping how the industry delivers value to its clients in the future.

It is noteworthy that the strength of Generative AI is in exploration and explanation rather than pure prediction, thereby, it could offer significant benefits in several key areas like real-time Q&A where seamless communication, and complex regulatory and compliance guidelines are readily explained through user interaction with Generative AI-powered chatbots. Through the integration of the GPT into an organization's communication channels, organizations can leverage Generative AI to interact with users, providing precise answers and valuable guidance, thereby enhancing organizational benefits. In personalized interaction where Generative AI can help tailor financial products and services to individual needs. Imagine a system that helps customers understand complex lending packages and allows account representatives to be more responsive to specific client situations. Furthermore, for increased efficiency and productivity, automated workflows where real-time transcription of calls with flagged indicators would allow for faster analysis and identification of areas requiring management attention. This promises to streamline workflows and improve decision-making. Generative AI can also help address the tech talent requirements and time criticality, Generative AI can automate repetitive tasks, freeing up human resources to focus on higher-value activities. This can significantly improve work processes pertaining to backlogs and improve overall operational efficiency.

Generative AI use cases extend further to specialized products which would offer users invaluable insights into intricate financial products, like facilitating understanding of lending packages. Moreover, in the realm of premium products, Generative AI could enhance existing services by ensuring greater accuracy, responsiveness, and the potential for a broader array of virtual advisors. Similarly, in high-end services such as wealth management and insurance, Generative AI serves as an intelligent assistant, complementing human expertise while ensuring human oversight of critical decision-making processes.

The impact of Gen AI integration is tangible and profound. A striking example is observed in the property and casualty (P&C) insurance sector, where Generative AI has significantly slashed claim cycle times from averaging 140 days to about 40 days[17]. This substantial reduction

[16] Morgan Stanley, "How Generative AI Could Reshape work", November 2023
[17] AI in Fintech - Bain & Company podcast interview with CJ Przybyl, Reserv Inc. November 2023

underscores the transformative potential of Generative AI in streamlining operations and delivering substantial efficiency gains and delivering value.

In summary, while the financial industry recognizes the potential use cases of Generative AI in improving internal processes and customer interactions, there are concerns over risk and customer trust. However, notable successes in advisory services, P&C claims, coding, and internal support, indicate that considerable productivity gains are achievable through strategic implementation of Generative AI. As companies continue to explore their capabilities and address challenges, the use cases for Generative AI will increasingly play a significant role in shaping the future of the financial industry. A crucial decision Financial institutions have to make is whether to build their own Generative AI capabilities or partner with existing firms providing solutions. It's important here to remember that Generative AI isn't a "plug-and-play" solution. Careful planning and integration are crucial for the successful integration of Generative AI.

Identifying the Numerator of ROI for Generative AI

Numerator Identification: The Role of COGS and SG&A

The effectiveness of Generative AI in a financial context is not only about deploying technology but also about its impact on the company's cost structure. The numerator in our ROI formula is crucial as it reflects the direct financial benefits obtained from the investment. Here, we consider Cost of Goods Sold (COGS) and Selling, General, and Administrative expenses (SG&A) as primary components. We derived the data from the financial statements of the companies.

To systematically evaluate the impact of Generative AI, we employ a two-by-two framework:

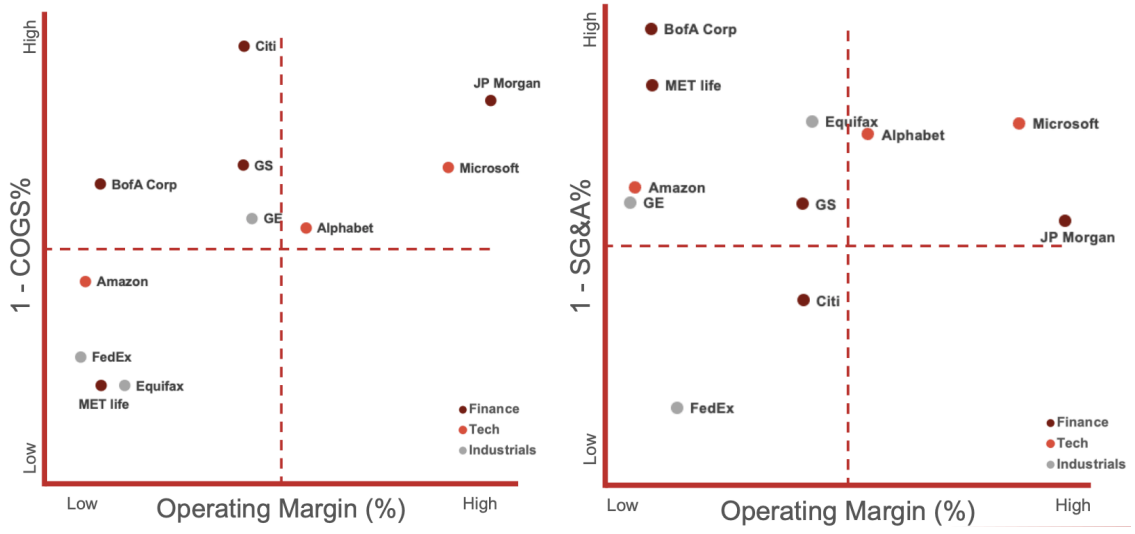

Chart 3: Two-by-two framework for COGS and SG&A

Y-Axis: Potential for Cost Structure Improvement
The Y-axis is defined as 1−COGS/Total Revenue. COGS typically includes direct costs like raw materials, labor, and manufacturing overhead. For some AI investments and in some industries, these direct costs are less relevant because AI primarily impacts operational efficiency and other indirect expenses. By focusing on the portion of total revenue not consumed by COGS, we are essentially highlighting the portion of the business that could be improved through AI. This includes areas like SG&A, which can often be optimized through automation, better decision-making, and efficiency improvements brought by AI technologies. Overall, for the Y-axis, the higher this value, the more room there is for a company to leverage AI to enhance its cost structure and operational efficiency.

X-Axis: Competitive Margin Improvement
The X-axis is defined as the Margin Percentage, calculated as Operating Income/Total Revenue. This measures how much of the revenue is retained as operating income after accounting for the costs that are relevant to the company's core operations but before financial expenses and taxes.

A higher margin percentage indicates that a company is able to generate more profit from its sales, potentially reflecting a competitive advantage. In the context of AI investment, an improvement in this margin suggests that the company is not just reducing costs but may also be enhancing its product or service offerings, leading to increased pricing power or market share in a competitive landscape.

The integration of Generative AI aims to enhance both these metrics by improving operational efficiency and competitive margins, leading to a substantial impact on the numerator of our ROI calculation.

Challenges
The primary challenge in using this framework is ensuring the accuracy and availability of data to measure these impacts. As Generative AI is a relatively new technological investment, historical data on its impacts are limited, which complicates empirical assessments.

Findings
Our analysis has shown that companies with high SG&A expenses, like those in the financial sector, have substantial room for efficiency improvements through Generative AI. This observation is particularly notable in companies like JP Morgan, where AI can significantly influence operational efficiencies and cost management.

The key findings from the table on the right indicate varied outcomes across different companies. JP Morgan, ranking first in the AI index, outperforms its peers in all aspects, demonstrating the positive impact of AI investment on its performance. However, the rise in SG&A per employee, primarily due to tech staff hiring and escalating talent costs, has increased pressure on the operating efficiency ratio in the short term, showing lower cost savings on SG&A per head.

Company *Unit: $million per person*	AI index and ranking[1]	Δ Revenue per Employee[2]	Δ SG&A per Employee[3]	Δ Operating efficiency[4]
JP Morgan	69.3 (1st) A	0.1% ▲	(0.9)% ▼	1.1% ▲▲
Capital One	62.4 (2nd)	(0.4)%	2.6% ▲ B	(2.4)%
Wells Fargo	42.2 (3rd)	(1.1)%	2.7% ▲	(4.1)%
Citigroup	37.5 (9th)	(2.5)%	(2.6)%	(2.7)%
BoA	35.3 (15th)	(0.7)%	(5.4)%	(2.6)%
Morgan Stanley	34.0 (17th)	(1.7)%	(2.7)%	(0.03)%
Prudential	NI	(0.9)%	(9.8)%	(3.1)%
Visa	NI	(4.5)%	(7.9)%	(3.1)%
Blackrock	NI	(2.8)%	(3.5)%	(0.7)%

Source: Income statement and employee data from https://stockanalysis.com/
AI index and ranking from https://evidentinsights.com/reports/key-findings-report?id=ae2ea186b8

1. AI index measures level of development, sophistication, and effectiveness of AI adoption.
2. Revenue is not macroeconomic adjusted.
3. Inflation and one-off expense adjusted.
4. Operating efficiency ratio = SG&A expense for the period divided by total net revenue for the period.

Chart 4: Analysis of select companies

In terms of revenue per employee, JP Morgan showed a marginal increase of 0.10% CAGR from 2014 to 2023. This suggests that while AI investment may lead to overall improved performance, it does not necessarily translate into significant revenue growth on a per-employee basis. In contrast, some companies like Citigroup and Visa, despite lower AI index rankings, showed substantial decreases in revenue per employee, indicating that other factors beyond AI investment might be influencing these outcomes over the past 10 years. It is important to note that the analysis does not isolate Generative AI-related impacts and is constrained by time frame limitations and other macroeconomic factors, such as business cycles that affect financial data and employee numbers.

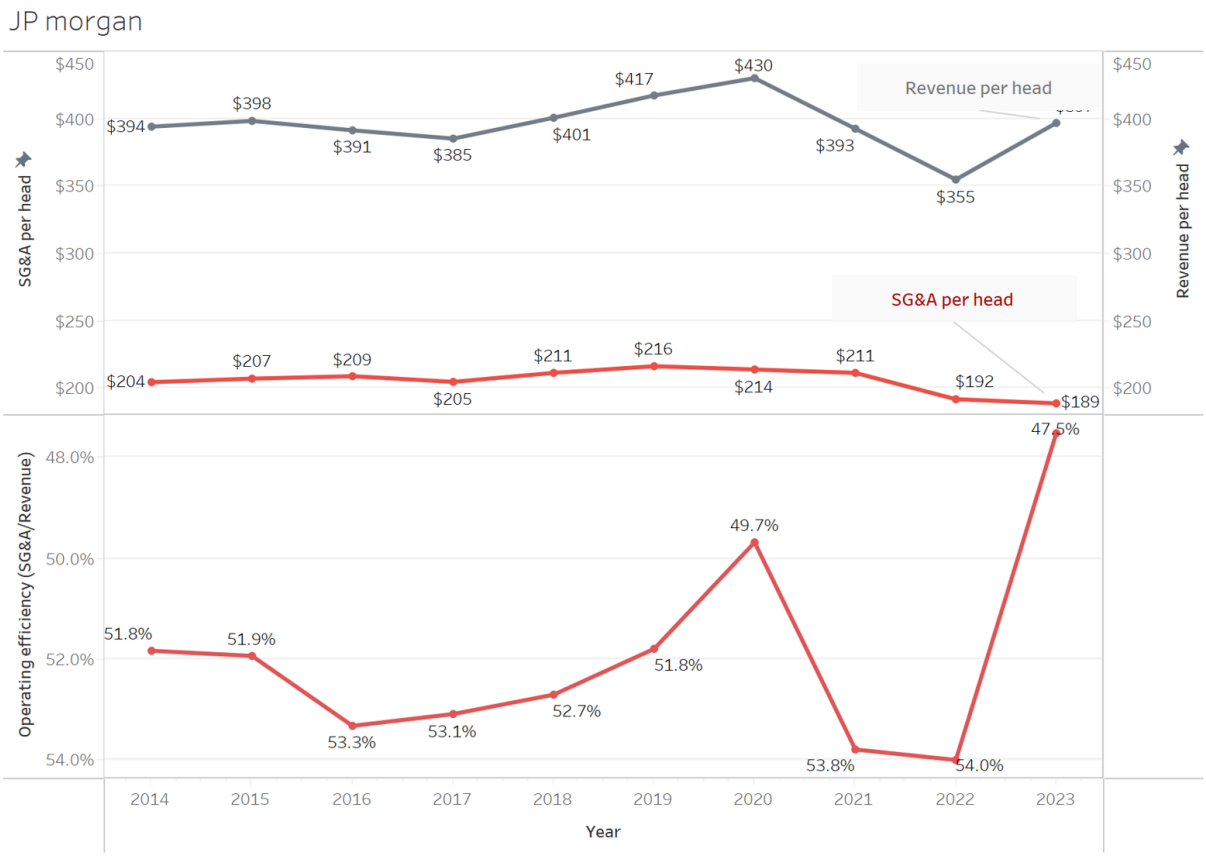

The trends of SG&A per head, Revenue per head and Operating efficiency (SG&A/Revenue) for Year Year. Color shows details about SG&A per head, Revenue per head and Operating efficiency (SG&A/Revenue). For pane Sum of SG&A per head: The marks are labeled by SG&A per head. For pane Sum of Revenue per head: The marks are labeled by Revenue per head. For pane Sum of Operating efficiency (SG&A/Revenue): The marks are labeled by Operating efficiency (SG&A/Revenue).

Chart 5: Year-on-year changes in JP Morgan's revenue per head, SG&A per head, and operating efficiency improvement

- **Intuition**
  Y-Axis: Better the product understanding, higher the potential to incorporate AI while keeping the costs manageable

  X-Axis: Higher the relevant data availability, AI can help improve the ease with optimization

- **Structure**
  - This is a broad distribution of products and industries over the matrix (Finance, Tech, and Industrials)

- **Conclusion**
  - Generative AI is could have higher impact potential, as a positive correlation as in the matrix
  - Thereby, setting up for a broad framework to analyze costs involved, the ROI on integrating Generative AI, could be suitably understood

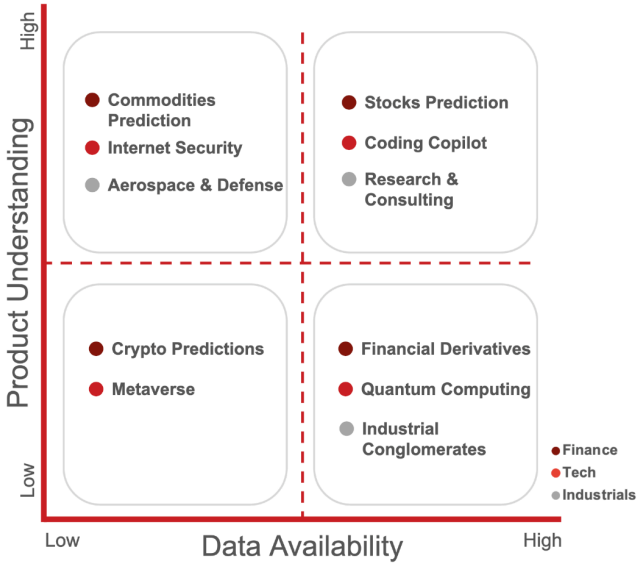

Chart 6: Product Understanding and Data Availability

To analyze different industries' and products' potential to integrate Generative AI, we put different products and industrial sectors into a matrix as shown in the chart above. This is to be understood in comparative terms around data availability and product understanding for the potential to integrate Generative AI. For instance, the relevant information about coding process is available which can help write codes, as well as the product, is well understood, therefore that figures in the top quadrant. Other products and industries are intuitively placed in the quadrants for comparative analysis per data availability and the product understanding in the context of integration potential for Generative AI.

Historical Approach

Both the Literature Review, and the approach to measuring the potential from Generative AI through COGS, and SG&A have come to the same conclusion. A lack of prolific and diverse data prevents us from moving forward.

The closest we have come to determining the ROI was through an incremental benefits approach. In principle, this is a rather common-practice backward time series analysis which would allow us to identify who has made investments and what benefits they will gain from it.

The idea was to compare an individual firm's metric (SG&A or COGS) before and after they implemented Generative AI, and then evaluate the difference in this metric to ascertain the generated return. Despite theoretically possible and common practice, the approach is practically not feasible at the moment. First and foremost, we are not able to distill the noise from SG&A or COGS. There are numerous factors that might cause an increase in either of the two metrics. Second, most firms have only recently implemented Generative AI. Thus, there is not enough time to use as a comparison. Finally, not many firms have implemented Generative AI, minimizing the pool of firms left to evaluate.

Once again, a widespread, straightforward, and logical approach to calculating the ROI proved to be unsuccessful. We now have come to the conclusion that pursuing analysis backward does not work as there is not enough data. Thus, finding a particular benchmark number based on historical data is not possible.

Lacking Data and Excess Noise

After conducting a historical assessment, we have determined that retrospective analysis is not feasible due to insufficient data and a lack of evidence to establish a causal relationship between historical data and future projections in this field. Only a small percentage of companies choose to disclose AI-related investments and earnings in the company's annual report, and most often do not disclose specific numbers or do not distinguish AI-related businesses separately. With this realization in mind, we embarked on exploring whether we can determine the viability of investing in a company that meets specific prerequisites. Therefore, rather than seeking a precise number, our focus should be on identifying a particular range.

Taking into account the considerations mentioned above, we opt to utilize the decision tree model for analysis. Decision trees serve as an effective tool for uncovering non-linear relationships between variables, particularly in cases where the existence of a causal relationship remains uncertain. This makes decision trees invaluable for predictive modeling tasks. Their utility stems from several advantages: they are easy to interpret, requiring minimal mathematical expertise to grasp their functioning; they exhibit robustness in handling missing values, rendering them suitable for real-world datasets plagued by incomplete information; moreover, decision trees excel in elucidating classification problems, providing clear insights into the underlying data patterns. However, despite these strengths, decision trees are not without their drawbacks. They are highly sensitive to random fluctuations in data, making them susceptible to the presence of 'white noise' that can distort their predictive accuracy. Additionally, decision trees employ a greedy algorithmic approach, which may lead to suboptimal solutions or local optima, particularly in complex datasets. Furthermore, explaining the decision process for continuous features can prove challenging due to the inherent discretization required by decision tree algorithms.

One of the main reasons we decided to forego further research was the limited availability of data. Despite our best efforts, we encountered significant challenges in obtaining sufficient data for analysis. The reason for this shortfall is that the novelty of the topic made it difficult to collect enough historical data to inform our decision-making process. When we use all the existing financial data obtained from annual reports to make predictions, the predictions of the constructed models are very poor, with an accuracy rate of usually less than 50%, and there is not enough data for model revision as well as testing. In addition, the inherent complexity of the topic poses an additional hurdle, especially in terms of identifying potential side effects associated with the product in question. Without a robust dataset containing a wide range of variables, it becomes extremely challenging to distinguish between causal relationships and spurious correlations, undermining the reliability of any predictive model derived from this limited data.

In addition, our decision to discontinue the product was due to the lack of a comprehensive understanding of the product itself. In order to gain a deeper understanding of the product's performance and potential impact, we explored the possibility of utilizing third-party services or developing our own predictive models. However, despite these efforts, we encountered significant barriers to gaining a deeper understanding of the product's complexity. The lack of insight into the product's functionality and its broader impact forced us to err on the side of caution and reevaluate our approach. If we had had a specific business line and application in mind, it may perhaps have facilitated our decision tree model at the risk of turning our project into a framework, so specific, that its applicability becomes limited.

Findings of Part I

In short, we can see that what should have been straightforward, theoretically, proved to be challenging and infeasible, practically. The conclusion that we have derived is the fact that lacking data and the necessity to become more specific hinder this approach from moving

forward. Regardless of which route there exists, modeling the ROI for Generative AI, generally, cannot be done at the moment due to the aforementioned reasons.

One way this framework could be pursued would be by looking into individual business functions and specific applications to determine the ROI of implementing Generative AI within tasks of these particular applications. But as mentioned above, this would reduce the applicability and relevance of the model and necessitates data that most likely only is available to the firm itself. Our analysis has showcased that due to lacking data and its divergence across applications, even within an industry, a holistic framework cannot work at this time.

PART II

The Social Benefit of Generative AI

The development and improvisations of Generative AI technology has allowed people and companies to unlock a new potential in their productivity. With the tool being able to generate works in various landscapes, from papers and reports to customer service responses, in a fraction of the time it used to take for people to complete, it is agreed that it has the potential to free the time it used to take humans to tackle the identical task. Companies seeking for a higher productivity gain and more efficient usage of their financial resources have shown great interest in the myriad of possibilities Generative AI can uncover and the new frontier that is facing them.

The time that is saved from assigning tasks to Generative AI instead of human employees is a great asset to companies. Companies can use this time to allocate more complex and value-adding tasks to their employees, aiming for higher quality creations, or save on their financial assets by simply substituting Generative AI tools in place of their human workers, and utilize the money on other investments that they believe would bring additional revenue to the firm. Whichever way, the time liberated or the money brought in, the saved asset can be seen as an increase in the opportunity cost of choosing not to deploy Generative AI in a company. As the perceived opportunity cost increases over time with the development of Generative AI and the elevated quality and quantity of its abilities, it is foreseen that firms will increasingly adopt the tool into their operations.

As a higher proportion of companies within a society deploy Generative AI, the collective revenue generated by and through these decisions can be translated into the value added to the society as a whole. In other words, the aggregated surplus revenue would be the value that is added to society by the technological advancements in Generative AI.

Society's increasing adoption of the technology will inevitably bring its costs down at a certain point in time, due to economies of scale. As costs are driven down, demand will become higher, and even more companies will be adopting the technology. Through the iterating cost fluctuations due to changes in demand and advancements in technology, Generative AI will become prevalent in society and will inevitably reach a point where no further value can be added to society through Generative AI. Companies will exhaust the value that can be added through deploying Generative AI, meaning that further investments in Generative AI will no more create revenue because all the tasks that can be done by Generative AI have already been taken. At this stage, Generative AI is effectively acting as a public good – similar to the internet.

The accumulated revenues that have been created by Generative AI from all companies within the society will be what the society benefits from adopting the tool. The next section further discusses how social welfare increases with the adoption of Generative AI as it alters the cost of goods.

The Social Utility Function

The effect that Generative AI would have on society can be explained using the social utility function. The social utility function represents the aggregate of the utilities of individuals given the goods and services within a society. Utilizing Generative AI is likely to yield an increase in individual utilities, through channels such as more efficient allocation of resources, more leisure time through the time saved by Generative AI, and for companies the lowered cost of achieving the same work. While this may not apply to every individual and some may experience a lowered individual utility by being adversely affected by the adoption of Generative AI, the society at the aggregate level will arguably experience a higher utility.

Once the additional utility brought by adopting Generative AI is exhausted and the social utility curve reaches a stable state, it will no longer enhance the social level of utility. This is because once costs have been sufficiently brought down and access to Generative AI tools has become normalized, the social utility curve will no longer be able to capture the effects of Generative AI on utility. This is because Generative AI at this stage will be acting similar to a public good[18] or merit good[19] and will not be able to be internalized within the model. In this long-run scenario, Generative AI is no longer a tool to create additional utility but a good that almost everyone in society has access to.

Advantages and Disadvantages of Generative AI

Socially, the use of Generative AI tools can liberate human labor from a multitude of tasks and lessen the burden for many people, generating and unlocking benefits that were out of reach previously. This is made possible due to the fact that first, Generative AI can automate repetitive tasks and allow for heightened efficiency and productivity in the workplace. There have been use cases where Generative AI tools such as ChatGPT have been used to automate repetitive tasks, such as customer service responses, to free up the time of employees to be allocated more towards comparatively complex and value-added tasks.[20] The use of Generative AI in customer service operations has been especially widely discussed, where there is ample evidence that customer experience can be improved when corporations deploy Generative AI tools in these functions.[21] Other fields that currently are receiving high expectations to be merited by Generative AI include education, both in the teaching, learning, and research aspects and healthcare.[22]

[18] Public goods are non-excludable and non-rivalrous goods, meaning that individuals cannot be effectively excluded from using them. Examples include clean air or street lighting.
[19] Merit goods are goods or services that the society believes everyone should have access to, such as education or sometimes electricity.
[20] Gupta, B.; Mufti, T.; Sohail, S.S.; Madsen, D.Ø. “ChatGPT: A Brief Narrative Review”. Cogent Business & Management 2023, 10.
[21] Ray, P.P. “ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope”. Internet of Things and Cyber-Physical Systems, Volume 3, 2023, 121-154.
[22] Bahrini, A.; Khamoshifar, M.; Abbasimehr, H.; Riggs, R. J.; Esmaeili, M.; Majdabadkohne, R. M.; Pasehvar, M. “ChatGPT: Applications, opportunities, and threats.” In 2023 Systems and Information Engineering Design Symposium, 274-279.

Additionally, Generative AI overcomes a critical loophole in human productivity, which is its accessibility regardless of time. This powerful characteristic of the tool uncovers a new area of possible working hours, as, unlike humans, Generative AI can be accessed around the clock at 24/7. Not only does this availability enhance the productivity of a company, but it also better addresses different users' needs across diverse areas around the world.

The downside to Generative AI, however, is that it is susceptible to generating biased and sometimes false information, based on the training data used. There have been many observed cases where statements generated from Generative AI tools were biased sexually or racially raising social concerns and eventually being terminated for use. False information can facilitate the spread of misinformation in areas where truthfulness is important, such as legal claims or advertising materials, and may cause significant problems to the entity that was using the AI tools.[23] This may imply additional costs incurred for legal consultations and for recovering from the damaged brand image, which may possibly overturn the benefits brought in by deploying Generative AI.

Other concerns when using Generative AI include those on privacy and cybersecurity attacks. The processes of data collection and training use a lot of data that is not always solicited by the data provider, and if used extensively can cause privacy problems also potentially leading to legal disputes. The applications of Generative AI in surveillance cause further privacy complications, implying that companies should pay careful attention to the coverage of Generative AI models when making utilization decisions.[24] Additionally, the increasing prevalence of cybersecurity attacks also introduces an additional layer to companies' responsibilities when using Generative AI.

Conclusion

In short, this paper has presented the merit of creating a framework to determine the Return on Investment (ROI) of Generative AI. As a formula for calculating the ROI on regular investments is rather simple and widespread, calculating the ROI for Generative AI should not have been too challenging.

We have first entertained a private benefit approach that analyzed numerous possibilities through which a formula could be designed, that would showcase monetary improvements associated with employing Generative AI. Specifically, we have analyzed the merit of pursuing an approach that measures the potential improvement through COGS or SG&A, pursuing a historical approach that uses prior data of firms that have already implemented Generative AI to associate that to a cost improvement to attain a specific number, or using a decision tree model to attain a range.

[23] Gupta et.al.

[24] Baldassarre, M.; Caivano, D.; Nieto, B.F.; Gigante, D.; Ragone, A. "The social impact of Generative AI: An Analysis on ChatGPT". GoodIT '23: Proceedings of the 2023 ACM Conference on Information Technology for Social Good. September 2023, 363-373.

We have come to the conclusion that calculating a framework to determine the ROI on Generative AI is not possible at the moment as not enough firms have implemented the technology yet, the technology is yet too nascent, and the associated data is scarce. Any actor attempting to create an ROI for Generative AI at this time will face the same challenges outlined in this paper and reach the same conclusion.

As aforementioned, one possible avenue to consider continuing with this project would be to look into individual business functions and specific applications to create a framework solely for a specific application. Risk and key metrics could easily be defined and associated data may be available internally to the firm. The culprit of this approach is that it would only apply to a specific application within a specific firm.

Finally, the social benefit of Generative AI has been evaluated, and a key takeaway would be that Generative AI will not take away people's jobs but will allow them to do things that they were not able to do previously. This would not only bring an overall benefit from the society's perspective but also create new jobs in return.

Bibliography

1. "The state of AI in 2023: Generative AI's breakout year." 2023. McKinsey & Company. August 1, 2023. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year#/.
2. PricewaterhouseCoopers. n.d. "Solving AI's ROI Problem. It's Not That Easy." PwC. https://www.pwc.com/us/en/tech-effect/ai-analytics/artificial-intelligence-roi.html.
3. Baxter, Kathy. 2023. "Managing the Risks of Generative AI." Harvard Business Review. October 24, 2023. https://hbr.org/2023/06/managing-the-risks-of-generative-ai.
4. Oliver Wyman, "How Generative AI Is Transforming Business and Society", 2024.
5. https://www.oliverwymanforum.com/content/dam/oliver-wyman/ow-forum/gcs/2023/AI-Report-2024-Davos.pdf
6. IBM, "Generating ROI with AI", 2023. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ai-capabilities
7. Chui, Michael, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel. 2023. "The Economic Potential of Generative AI: The Next Productivity Frontier." *McKinsey & Company*. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#introduction.
8. "How Generative AI Can Boost Highly Skilled Workers' Productivity | MIT Sloan." 2023. MIT Sloan. October 19, 2023. https://mitsloan.mit.edu/ideas-made-to-matter/how-generative-ai-can-boost-highly-skilled-workers-productivity#:~:text=A%20new%20study%20on%20the,who%20don%27t%20use%20it
9. "AI Investors: Act Fast, Act Wisely." 2023. Bain. October 19, 2023. https://www.bain.com/insights/ai-investors-act-fast-act-wisely-tech-report-2023/.
10. Raimondo, Gina M., Laurie E. Locascio, U.S. Department of Commerce, and National Institute of Standards and Technology. 2023. "Artificial Intelligence Risk Management Framework (AI RMF 1.0)." Report. *NIST AI 100-1*. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.
11. Morgan Stanley, "How Generative AI Could Reshape work", November 2023 https://www.morganstanley.com/ideas/generative-ai-future-of-work
12. Bain & Company - "AI in Fintech" - podcast interview with CJ Przybyl, Reserv Inc. November 2023 https://open.spotify.com/episode/4joTzln2Xk7PE7dNOqVd8K?si=779f17b40d9e49e7
13. Gupta, B.; Mufti, T.; Sohail, S.S.; Madsen, D.Ø. "ChatGPT: A Brief Narrative Review". Cogent Business & Management 2023, 10.

14. Ray, P.P. “ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope”. Internet of Things and Cyber-Physical Systems, Volume 3, 2023, 121-154.
15. Bahrini, A.; Khamoshifar, M.; Abbasimehr, H.; Riggs, R. J.; Esmaeili, M.; Majdabadkohne, R. M.; Pasehvar, M. “ChatGPT: Applications, opportunities, and threats.” In 2023 Systems and Information Engineering Design Symposium, 274-279.
16. Baldassarre, M.; Caivano, D.; Nieto, B.F.; Gigante, D.; Ragone, A. “The social impact of Generative AI: An Analysis on ChatGPT”. GoodIT '23: Proceedings of the 2023 ACM Conference on Information Technology for Social Good. September 2023, 363-373.