

MOT2910 Master Thesis Project

Overcoming the Barriers to Large Language Model (LLM) Adoption: A study on Organisations' Perceived Risks of LLMs

Delft University of Technology
Technology, Policy & Management

M.J.M. Kooij
5318467

Chair:	Dr. H.G. van der Voort, Section Organization & Governance
First Supervisor:	Dr. Y. Zhauniarovich, Section Organization & Governance
Second Supervisor:	Dr. M.B.O.T. Klenk, Section Ethics/Philosophy of Technology
External Supervisors:	V.S. de Azevedo, MSc, Cybersecurity Consultant Ir. B.A.P. van den Kieboom, Cybersecurity Consultant



Contents

Executive Summary	1
Nomenclature	4
List of Figures	5
List of Tables	5
1 Introduction	6
1.1 Relevance	6
1.1.1 Academic Relevance	6
1.1.2 Societal Relevance	6
1.1.3 MoT Perspective	7
2 Literature Review	8
2.1 Relevant Terminology	8
2.1.1 GenAI vs LLM	8
2.1.2 Knowledge Work	8
2.2 GenAI in the Workplace	9
2.2.1 LLM Adoption and the Impact on Work	9
2.2.2 Employee Perceptions of LLM Benefits for Work	9
2.2.3 Effects on Employee Productivity and Quality	10
2.2.4 Organisational Policies on GenAI Use	11
2.3 Risks of LLMs for Business	12
2.3.1 Identifying Risks of GenAI Development	12
2.3.2 Job Displacement due to Automation	12
2.3.3 Cybersecurity Risks	13
2.3.4 Ethical Risks and Blind Spots	14
2.3.5 Non-Compliance with AI Regulatory Requirements	14
3 Research Methods	16
3.1 Research Gap	16
3.2 Research Objective	16
3.3 Research Scope	16
3.4 Research Questions	16
4 Methodology	18
4.1 Qualitative Research	18
4.1.1 Technology Acceptance Model	18
4.2 Participants	19
4.3 Data Collection & Processing	20
4.3.1 Interview Structure	20
4.3.2 Interview Transcription	21
4.3.3 Transcript Coding & Grouping	21
4.4 Ethical Considerations	22
5 Analysis of LLM Adoption Perceptions	23
5.1 Benefits Overview	23

5.2	Risks Overview	24
5.3	Current LLM Adoption	26
5.3.1	Authorization of LLMs	26
5.3.2	Usage Policies	28
5.3.3	Employee LLM Usage & Most Common Use Cases	29
5.4	Future LLM Adoption	32
5.4.1	Future LLM Adoption Capabilities	33
5.4.2	Adoption Decision Process and Strategies	33
5.5	Employee Perspective	34
5.5.1	Employee Perceived Benefits	34
5.5.2	Employee Perceived Risks	35
5.6	Organisational Perspective	36
5.6.1	Organisational Perceived Benefits	36
5.6.2	Organisational Perceived Risks	36
5.7	Employee Expectations of LLMs	37
6	Discussion	39
6.1	Technology Acceptance Model for LLM Adoption	39
6.2	Validity & Reliability	41
6.3	Study Limitations	41
7	Conclusions	43
7.1	Results Discussion	43
7.1.1	Most Common LLM Use Cases	43
7.1.2	Employees' Perceived Benefits and Risks of LLM	43
7.1.3	Organisations' Perceived Benefits and Risks of LLM	44
7.1.4	Future Adoption and Employee Expectations	44
7.2	Research Conclusions	45
7.3	Future Work	46
	Appendix A: Employee and Organisational Perceived Benefits and Risks of LLM Adoption	47
	Appendix B: Summary of Future LLM Adoption Plans	48
	Bibliography	49

Executive Summary

Rapid advancements in Generative Artificial Intelligence (GenAI) have developed the capabilities to produce content that is increasingly indistinguishable from human-generated work. This trend was especially marked by the release of OpenAI's ChatGPT in late November 2022 (OpenAI, 2022), which is one of several Large Language Models (LLMs) widely available for public use (Dell'Acqua et al., 2023). GenAI is expected to have a significant impact on how business is done, especially for knowledge-intensive domains. Professional usage amongst knowledge workers is already widespread and many of them believe that LLM use for work will make them more efficient, help them generate ideas, and improve the quality of their work. Yet, the technology comes with numerous potential risks, including job displacement, threats to data privacy, unreliability, cybersecurity risks, and non-compliance with new AI regulation. The decision to adopt LLMs for knowledge work is thus a significant one as organisations must carefully weigh up the benefits and drawbacks that these models present.

Using a qualitative research approach, this study was aimed at exploring employee and organisational perceptions on the benefits and risks of LLM adoption within the Dutch financial sector. This study was conducted in partnership with a global Professional Services Firm and used their existing network of people and clients to conduct interviews with experts who act as advisors to top management in the decision-making process of new technology adoption such as LLMs. 18 semi-structured interviews were done to collect qualitative data. Purposive sampling was used for selecting interview participants, and thematic analysis of interview transcripts was used in the data analysis. The findings were then adapted to the Technology Acceptance Model (TAM) to draw conclusions on how organisations should best handle LLM adoption. The research question to be answered is:

- How do employees' and organisations' perceived benefits and risks of Large Language Models (LLMs) influence financial organisations' LLM adoption plans?

To get to the research question's answer, the following sub-questions were explored first.

- SQ1. What are the most common use cases of LLMs?
- SQ2. What are employees' perceived benefits and risks associated with LLM adoption?
- SQ3. What are organisations' perceived benefits and risks associated with LLM adoption?
- SQ4. What are organisations' current levels of LLM adoption and their future LLM adoption plans, and how are they aligned with employee expectations of LLMs?

The five overall most common perceived benefits for the adoption of LLMs are efficiency/productivity gains (16%), the automation of repetitive tasks (7%), facilitating information search (5%), improving writing/grammar and translation (5%), and customer service improvement (5%). Some of the least mentioned benefits (0.4%) include scalability, consistency, versatility, and an increased understanding of LLMs. The five overall most common perceived risks for the adoption of LLMs are the lack of quality control/output validation (10%), inputting or exposing sensitive data (9%), the risk of a data leakage/breach (7%), job loss due to LLM automation (6%), and user manipulation (6%). Some of the least commonly mentioned risks include LLM training data poisoning, social engineering, single point-of-failure, and reduced job satisfaction (0.3% each).

Exactly half of participants (9/18) said that they use LLMs in their own work while the other half stated that they do not due to a lack of perceived benefits. Microsoft was found to be a significant player in the current adoption of LLMs at Dutch financial institutions with 57% of all LLMs used by participants being owned by the software vendor. The most common LLM use cases were literary and creative in nature and included preparing presentation slides (16%), text generation (12%), email composition (12%), and structuring documents (12%).

On organisations' current levels of LLM adoption, it was found that only a minority of financial organisations prohibited employees from using LLMs for their work due to not having a good enough understanding of the technology, not having the proper safeguards in place, the risk of exposing sensitive data that is inputted in the LLM, and the organisation's lower risk appetite. 78% of participants said that their organisation allows LLMs at work, 17% of which allow it with limited use, and 22% said that they are prohibited. In terms of future LLM adoption plans, 72% of participants were aware of their organisation's plans, 17% were unclear, and 11% said that their organisation does not (yet) have plans. Most organisations (78%) have already taken the first step of adoption by running one or several LLM pilots to test the technology and its usefulness.

Current usage policies can be categorised in three ways: a complete ban on LLMs to eliminate the risks of LLM usage, or at least temporarily until new LLM policies are finalised and implemented; a partial ban which is the most common type, whereby certain LLMs are prohibited but not others, or whereby LLMs are allowed to select roles/departments or use cases (e.g., pilot development); and free use which is the least common type, whereby firms trust their employees to use LLMs responsibly and rely on their common sense. Most organisations have also already implemented pilots to assess future use cases of LLMs prior to adoption and research their added value. An emphasis is placed on risk assessments to make sure the risks are acceptable before continuing and so that the necessary safeguards can be implemented.

There is pressure on organisations to adopt quickly due to competition, with one firm taking a 'Smart Follower' approach which lets competitors make the first moves to reduce adoption risk while remaining one step behind to keep up with competitors. Meanwhile, another firm is under pressure to adopt as the lack of access to LLMs for programmers is creating a culture issue. Programmers are resorting to using workarounds to access these tools for their work and some are leaving the company for positions at competing firms that do have LLMs available. Barriers to future adoption plans include the lack of usage policies in place, and especially for pension funds, regulatory scrutiny, and a lower risk appetite due to the long-term nature of their investments. Finally, the decision to license or develop LLMs is an iterative process for which firms should consider the costs of ownership and maintenance, as well as the impact on business vs the effort to produce.

In terms of future LLM adoption plans, the targeted capabilities that organisations would like to achieve with their future adoption plans include helping programmers write better code, analysing help desk conversations with speech-to-text and gathering new insights, assisting employees via LLM chatbots, analysing emails to predict customer questions, and querying LLMs to return company documents from a database. Employees were found to be most excited about the potential efficiency and productivity gains that LLMs offer for their work (19%), how LLM usage could free up more time for focused, interesting, and fun work (11%), repetitive tasks becoming automated by LLM (8%), the new opportunities that LLMs present (8%) such as new business models, and improvements to customer service (5%). Comparing these two sets, general overlap can be seen between the targeted capabilities of future LLMs to be adopted and employee expectations, such as the productivity gains from helping programmers write better code or improvements to customer service as from using GPT to analyse emails and predict customer questions.

The five most common perceived employee benefits when it comes to using LLMs for work are increased efficiency/productivity (11%), improved writing/grammar/translation (9%), easier information search (9%), assistance with structuring documents (9%), and automating repetitive tasks (7%). The five most common perceived employee risks when it comes to using LLMs for work are inputting or exposing sensitive data via LLMs (14%), a lack of quality control or validation of LLM outputs (12%), developing an overreliance on LLMs (12%), model unreliability (10%), and bias present in LLM output or training data (6%).

The five most common perceived organisational benefits when it comes to using LLMs for work are increased efficiency/productivity (28%), automating repetitive tasks (13%), gaining a competitive advantage (6%), cost savings (6%), and customer service improvement (6%). The five most common perceived employee risks when it comes to using LLMs for work are inputting or exposing sensitive data via LLMs (14%), a lack of quality

control or validation of LLM outputs (12%), developing an overreliance on LLMs (12%), model unreliability (10%), and bias present in LLM output or training data (6%).

Finally, the Technology Acceptance Model was used as a theoretical framework to bring together the interview findings. Applying the model to LLM adoption showed that there is both a high perceived ease of use as LLMs are often interfaced through chatbots in natural language, and a high perceived usefulness as they improve employees' ability to achieve their most common work tasks quicker and more efficiently. This perceived usefulness of LLMs positively affects employee attitudes towards using LLMs. However, there are a few important risks perceived by organizations, notably risks of sensitive data being leaked via LLMs and a lack of quality control or output validation by employees using model outputs in their work. These risks seem to be disproportionately concerning to organisations who risk reputational damage while employees using the LLM are more focused on the efficiency gains offered by these models for their work tasks. This presents a misalignment of the employee and organisational perspectives.

Instead of restricting LLM usage, it is recommended that organisations find ways to incorporate them into their employee workflows by providing clear policies and guidelines. It was found that this approach will be most beneficial to creative and literary workflows like improving writing/grammar/translation, information search, structuring documents, and text generation. To achieve this integration, organisations should write and implement clear usage policies. This will ensure that the benefits of allowing LLM usage at work are enjoyed while mitigating the most important risks. Clear communication of items like which uses are and are not allowed, who is allowed to use it, and what kind of data may be inputted is crucial to avoid confusion. These specifications can be tailored to the needs of the organisation and according to their risk appetite.

As LLM is a relatively new technology, a limitation of the study is that it cannot be assumed that all participants have a good understanding of what they are and how they work. Moreover, policies and regulations for LLMs are still being written or have only recently been published. Similarly, AI roles are somewhat new in organisations, and there is no widespread set of AI roles expected for all organisations.

Nomenclature

Abbreviation	Definition
AI	Artificial Intelligence
API	Application Programming Interface
GenAI	Generative Artificial Intelligence
GPT	Generative Pre-trained Transformer
HREC	Human Research Ethics Committee
LLM	Large Language Model
MoT	MSc Management of Technology
MS	Microsoft
OWASP	Open Worldwide Application Security Project
TAM	Technology Acceptance Model

List of Figures

Figure 1: LLM is a subset of GenAI, which is in turn a subset of AI.	8
Figure 2: In the Jagged AI Frontier, tasks with the same perceived difficulty may be on one side or the other of the frontier (Ibid.).	11
Figure 3: Technology Acceptance Model flowchart. (Demystifying the Technology Acceptance Model, 2024) ..	18

List of Tables

Table 1: OWASP Top 10 risks and descriptions (Wilson & Dawson, 2023).....	13
Table 2: Interview participant overview including organisation, role, experience, and highest education.	20
Table 3: Organisation count per number of participants and unique organisations.	20
Table 4: Thematised benefit codes ordered by frequency, including percentage share of total benefit codes. .	23
Table 5: Frequency analysis of ‘benefits’ themes, ordered by total frequency of codes per theme.	24
Table 6: Thematised risk codes ordered by most frequent, including percentage share of total risk codes.	25
Table 7: Frequency analysis of ‘risks’ themes, ordered by total frequency of codes per theme.	26
Table 8: LLM authorization per organisation (Q5).....	27
Table 9: Summary of LLM usage policies per organisation.	28
Table 10: LLMs used by participants and LLM use cases.	30
Table 11: Participant LLM usage, ordered by most frequently mentioned (Q4a).	30
Table 12: Participant LLM use cases, ordered by most frequently mentioned (Q4b).	31
Table 13: Participant LLM use case themes, including number of codes per theme and total frequency.	31
Table 14: Participant responses on future LLM adoption plans (Q6), including mentions of active LLM pilots. ...	32
Table 15: Summary of future LLM adoption plans per organisation.	32
Table 16: Perceived employee benefits of LLM adoption for work.	34
Table 17: Perceived employee risks of LLM adoption for work.....	35
Table 18: Perceived organisational benefits of LLM adoption for work.	36
Table 19: Perceived organisational risks of LLM adoption for work.	37
Table 20: Participant responses on what they are most excited about regarding LLM adoption (Q13).	38
Table 21: Perceived benefits of LLM that participants are most excited about, ordered by frequency.	38
Table 22: Themes of perceived benefits that participants are most excited about, ordered by frequency.	39
Table 23: Interview responses (coded) on employee and organisational perceived benefits/risks (Q7 -Q10).....	47
Table 24: Summary of LLM future adoption plans including organizations’ use of pilots.	48

1 Introduction

In recent years, rapid advancements in Generative Artificial Intelligence (GenAI) have developed the capabilities to produce content that is increasingly indistinguishable from human-generated work. This trend was especially marked by the release of OpenAI's ChatGPT in late November 2022 (OpenAI, 2022), which is one of several Large Language Models (LLMs) widely available for public use (Dell'Acqua et al., 2023).

GenAI is expected to have a significant impact on how business is done, especially for knowledge-intensive domains. Organisations are thus looking to understand how this new technology should be adopted into their activities. From the figures, it is clear that businesses plan to adopt the technology, with Goldman Sachs forecasting global investment in GenAI to approach US\$200b by 2025 (Goldman Sachs, 2023). Yet, a 2022 IDC survey found that only 22% of organisations reported AI to be implemented on a large scale as part of the enterprise (Diasio, 2023).

One factor holding organisations back from LLM adoption may be the various perceived risks associated with such a new technology. Some examples include reliability, such as the accuracy of results and the potential for bias; intellectual property concerns; data privacy and security, especially for the safeguarding of sensitive information and trade secrets; and regulatory compliance. However, the current understanding of the perceived risks around LLMs that affect enterprise adoption is still limited. There is a need to investigate the various factors that influence adoption, including the most common use cases, employee and organisational perceived benefits/risks, and current adoption strategies and timelines, and how these are aligned with what employees and organisations expect from LLMs. These insights will hopefully inform organisations on how to improve their future adoption strategies.

1.1 Relevance

This section discusses the academic and societal relevance of this study, as well as how it fits within the MSc Management of Technology (MoT) study programme.

1.1.1 Academic Relevance

Organisations must constantly strive to improve their products, services, and business models to remain competitive. As such, most companies need to continuously acquire new technologies and apply strategic management and effective decision-making to innovate their business. This research explores the strategic management of LLM adoption. Specifically, there is a focus to uncover which of the risks and benefits of LLMs known in the existing literature are most of concern to employees and organisations in practice, to shed light on the way this new technology is being adopted within industry. This research is especially academically relevant as it explores the adoption of a relatively novel technology which has not yet been extensively studied. This fact raises questions related to the appearance of new technologies, such as whether they are a net positive to societies, which dangers should be addressed, and how the technology can be used to improve people's lives. The answers to these questions are crucial both for the researchers who are charged with the further development of LLM as a technology, as well as academics who advise on the formulation of new GenAI regulations. The open-ended approach taken during this study allows for a broader capture of perspectives and paves the way for future research to explore in more depth how particular factors influence the LLM adoption decision-making process.

1.1.2 Societal Relevance

There is a societal relevance to both organisations and society for a study on the perceived risks and benefits of LLMs and how they influence LLM adoption. Organisations want to know about the most important risks and benefits of the technology to their business, what kind of impact these will have on their operations, and how they can keep up with the competition. The manner and pace of LLM adoption therefore has importance to business success, innovation, productivity, and the wider economy. Beyond the direct interests of organisations, LLMs also have an influence on multiple facets of people's lives. The manner and pace of LLM

adoption within business can therefore have significant consequences for how people work, employee job security, and how customers interact with businesses, but also for the way personal data gets used and a person's ability to distinguish AI from human-generated content. Overall, the way LLMs are adopted by businesses will influence public perception and trust in the technology and in the organizations deploying them, which itself is an essential part of its widespread acceptance and adoption within society.

1.1.3 MoT Perspective

The MSc Management of Technology programme focuses on solving problems that are socio-technical in nature, where technology and people are intertwined. Students are taught to analyse technology, their commercial and societal impact, and implement these in the organisational context of a firm. This involves engineers “investigating and understanding, both internal to their own organisation and external in relation with business partners, what the current and future technological, economic, and social environments require technological firms to do” (*MSc Management of Technology, 2024*). In this research project, LLM is the technology being analysed and the objective is to uncover its most relevant perceived risks and benefits to understand how this technology can best be implemented in the organisational context of financial institutions. In doing so, adoption plans can be more thoughtfully designed, leveraging the technology's strengths, and managing its (perceived) risks, while addressing any misaligned perspectives between employees and the organisation. The research problem is socio-technical in nature as it focuses on the human perceptions of a new technology.

2 Literature Review

Due to their relatively recent emergence, LLMs have had fewer opportunities to be studied as compared to earlier AI technology. Although this limits the pool of existing literature on the subject, it in turn presents a myriad of new research opportunities. This section reviews the existing literature on LLMs within a professional work context, separated into three themes: relevant terminology, GenAI in the workplace, and the risks of LLMs.

2.1 Relevant Terminology

Before a summary of the literature can be made, it is important to understand the meaning and nuances of a few key terms. This section explores the nuance between GenAI and LLM and describes what is understood by the 'knowledge work' that LLMs are used for. These terms are relevant to the scope and context of this research when designing the study (e.g., targeting knowledge workers for interview).

2.1.1 GenAI vs LLM

GenAI and LLM are both terms within the field of AI. Whereas the former is a broader term referring to AI technologies capable of generating novel content, LLM is a specific type of GenAI model, specialized in understanding and generating human language. Although all LLMs are GenAI, not all GenAI models are LLMs (Figure 1) as some GenAI models generate entirely non-linguistic outputs (e.g., images, music).

$$\text{LLM} \subset \text{GenAI} \subset \text{AI}$$

Figure 1: LLM is a subset of GenAI, which is in turn a subset of AI.

Unlike most AI models, GenAI models are further distinct because they generate predictions by using the patterns of the data they were trained on, rather than relying solely on rules programmed into them. It is thus the generalised application of GenAI models which makes them stand apart from the traditional, procedural approach of earlier AI models (Orchard & Tasiemski, 2023).

2.1.2 Knowledge Work

As AI and human capabilities increasingly overlap with each other, the integration of human work with AI poses new fundamental challenges and opportunities. The focus of this study is on knowledge work given that the emergence of LLMs and their continued development bring about a novel and increasingly large threat to workers in knowledge-intensive domains. Earlier forms of AI had technical limitations which made it difficult to codify non-routine tasks. This made these tasks seem protected from automation, especially as previous waves of technology had mostly automated lower-skilled occupations. With the advent of LLMs, however, these models proved unexpectedly capable at creative, analytical, and writing tasks, representing an entirely new category of automation whose abilities overlapped with knowledge work. (Dell'Acqua et al., 2023)

Various definitions of knowledge work can be found in publications dating as far back as 1962, with certain themes having become increasingly common, such as a high level of education and skills, and the use of information technology as an integral part of the informational labour process (Pyöriä, 2005). One interpretation refers to organisational activities and occupations that are "characterized by an emphasis on theoretical knowledge, creativity, and use of analytical and social skills (Frenkel et al., 1995, p. 773)." In this kind of work, knowledge acts as the main input, the major way of achieving the work, and the output itself. Knowledge workers are then those whose major work tasks involve the creation of new knowledge or the application of existing knowledge in new ways. They typically have high levels of education and specialist skills, enabling them to identify and solve problems, and are the organisation's primary means of production. (Newell et al., 2009, p. 24)

With this clarification of the relevant terminology, the next sections explore the literature on GenAI within the workplace.

2.2 GenAI in the Workplace

This section brings together a variety of literature on GenAI as related to the workplace, including how LLMs will have a larger impact on knowledge work than non-GenAI models, employee views on the benefits of LLMs for knowledge work, LLMs' effects on worker productivity and quality of work, and employee views on (and awareness of) LLM usage policies within organisations.

2.2.1 LLM Adoption and the Impact on Work

As more and more organisations adopt LLMs for professional use, the technology is expected to transform the way in which work is performed. The following signals show that LLMs may change how most knowledge workers work. Firstly, major software vendors have already started integrating these technologies into all their core products (Cardon et al., 2023). Microsoft, for example, is working on integrating Copilot (based on the GPT-4 LLM) across its products like Bing, Edge, Microsoft 365, and Windows (Spataro, 2023). Secondly, the rate of LLM adoption by the public has been unprecedented, showing a strong appeal and perceived benefits of the technologies (Cardon et al., 2023). ChatGPT shocked the world when it set a new record for the fastest-growing user base, reaching 100m monthly active users just two months after launch (Hu, 2023). Third, new GenAI use cases are constantly being documented for many types of work (Cardon et al., 2023).

The generative capabilities of newer AI models like LLMs are furthermore expected to have a much more rapid and widespread impact on knowledge work than previous AI models. Three aspects of LLMs are suggested to explain this greater impact: the specialist knowledge provided by LLMs, worker performance improvements made easily accessible, and the "relative opacity" of LLMs. Firstly, although they are trained as general models, LLMs nonetheless demonstrate specialist knowledge and capabilities which are novel and unexpected, widely applicable, and quickly increasing. The second aspect is their ability to directly increase the performance of workers who use these systems, without the need for substantial organisational or technological investment. The third aspect is called "relative opacity", referring to LLMs' unclear failure points. Observed examples include the tendency to produce incorrect but plausible results, as well as a difficulty to predict how a model achieves good performance in some tasks but fails in others (known as the "jagged AI frontier", as seen in Figure 2). A lack of clear guidelines provided by developers on the best ways to use the models further contributes to this "relative opacity," leading users to resort to trial-and-error and knowledge sharing online. (Dell'Acqua et al., 2023)

This expected transformation of the workplace will also impact employees and the way they work, affecting their productivity and necessitating new policies to clarify how LLMs are to be used. These topics are explored further in the next sections.

2.2.2 Employee Perceptions of LLM Benefits for Work

Employee perceptions surrounding a new technology play a relevant and influential role in an organisation's decision-making process on the adoption of a new technology. If employees are enthusiastic about the technology and perceive it to be beneficial for their work, the organization will be more motivated and have more pressure to adopt the technology quickly (Davis, 1987). On the other hand, should employees not see the value of a new technology or view it to be detrimental to performing their work, the organisation will have less incentive to adopt it as support for adoption will be lesser and the likelihood of an inefficient allocation of resources greater. This section explores the literature on employee perceived benefits of LLM.

The paper "Generative AI in the Workplace: Employee Perspectives of ChatGPT Benefits and Organisational Policies" by Cardon et al. (2023) studies how GenAI affects employees in their research and communication tasks. Two studies were conducted, comparing early versus non-users of ChatGPT and employees of varying managerial status: non-managerial, managers, and executives.

The first study sought to find out the attitudes of professionals on the impact of AI on society and their jobs. About half of respondents thought that AI is good for society (43%), that it will do more good than harm (46%), and that it will increase productivity (52%). On the other hand, views varied widely based on how much experience participants had using ChatGPT. Those who had used it more than five times, for instance, were about three times as likely to think that AI will help them in their jobs compared to non-users. (Ibid., 2023)

The second study explored the different ways in which professionals use ChatGPT, what they perceived as the benefits of GenAI, and the perceived benefits of organisational policy surrounding its use. The results showed that ChatGPT usage among professionals is widespread: 42% have used it to research a topic, 32% to draft an email or text, 26% to draft text for a longer document (e.g., report), 21% to edit text, and 22% to summarize text. The study found that executives and managers appear to be using it more than non-managerial workers, especially for research and longer documents, and are slightly more likely to be enthusiastic about the benefits of GenAI. What is also interesting is that 71% of executives believe it can make them more efficient, as compared to just half of managers and non-managerial staff thinking the same. (Ibid.)

In summary, perceptions on employee perceived benefits of LLM usage appears to vary widely based on users' own level of adoption, ChatGPT usage among professionals was already widespread in March 2023, and higher status workers seem to employ the tool the most often. Large majorities of respondents, especially early adopters, believe GenAI will make their work more efficient, help them generate ideas for work, improve the quality of their work, and support more effective communication. Still, it is not yet known in which way these views influence organisational adoption of LLMs.

2.2.3 Effects on Employee Productivity and Quality

A significant expected benefit of LLM adoption is the increase in work productivity and quality that these tools offer. In an experiment involving 758 consultants from the global management consulting firm Boston Consulting Group, AI performance implications were examined on realistic, complex, and knowledge-intensive tasks with subjects randomly given one of three conditions: no AI access, GPT-4 AI access, or GPT-4 AI access with a prompt engineering overview.

The study proposes the idea of a "jagged AI frontier." It is observed that some tasks (like idea generation) are easy, while other tasks that seem easy (like basic math) are challenges for some LLMs. This is explained by a jagged frontier (depicted in Figure 2), where tasks that appear to be of similar difficulty may be performed better or worse by humans using AI. The "jagged" nature of the frontier means that the same knowledge workflow of tasks can have tasks on both sides of the frontier: tasks within the frontier are easily completed by AI, while those outside are beyond AI's current capability. (Dell'Acqua et al., 2023)

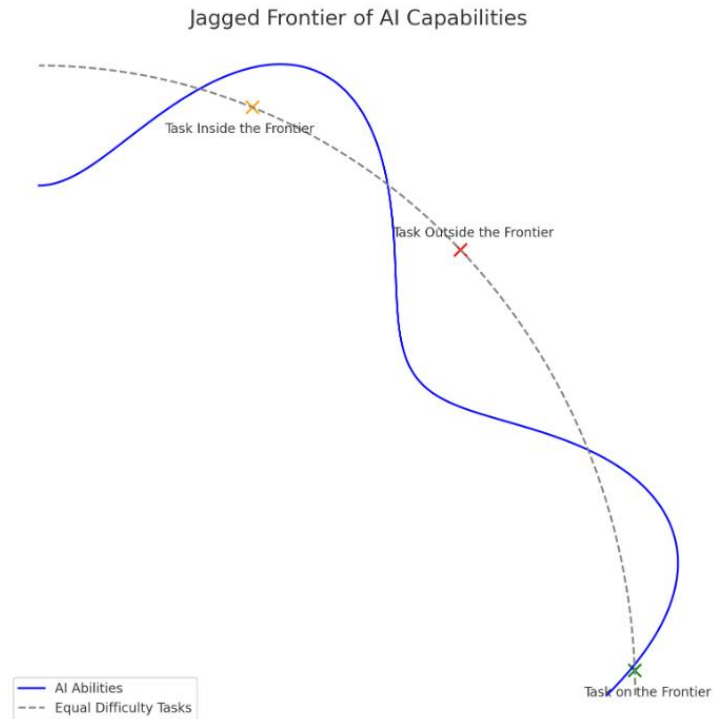


Figure 2: In the Jagged AI Frontier, tasks with the same perceived difficulty may be on one side or the other of the frontier (Ibid.).

For each one of a set of 18 realistic consulting tasks within the frontier of AI capabilities, consultants using AI were significantly more productive (completing 12.2% more tasks on average, and completing tasks 25.1% more quickly), and produced significantly higher quality results (more than 40% higher quality compared to a control group). Consultants benefited significantly from having AI augmentation, with those below the average performance threshold increasing by 43% and those above increasing by 17% compared to their own scores. Interestingly, this result implies that LLM tools like GPT-4 benefit below average workers more within the frontier. For tasks outside the frontier, however, consultants using AI were 19 percentage points less likely to produce correct solutions compared to those without AI. (Ibid.)

In short, AI can complement or even displace human work within the “jagged technological frontier”; outside of the frontier, AI output is inaccurate, less useful, and degrades human performance. This frontier is growing as AI capabilities rapidly evolve. Given these capabilities are often poorly understood, it can be hard for professionals to grasp exactly what the boundary of this frontier might be at a given time. (Ibid.)

2.2.4 Organisational Policies on GenAI Use

While the use of LLMs like ChatGPT have seen a swift rise in popularity, many organisations have struggled to clearly define policies and regulations surrounding its use in the workplace. A study from late January 2023 found that 43% of professionals have used AI tools, including ChatGPT, for work-related tasks, with nearly 70% of those professionals doing so without their boss’ knowledge (Graham, 2023). Only a few months later, large organisations like Apple, Samsung, and Amazon, as well as many banks like Goldman Sachs and Deutsche Bank, started banning the use of public LLMs in their workplaces over privacy concerns about the handling of sensitive company data and for regulatory reasons (Ray, 2023).

In Cardon et al.'s (2023) study on the perceived benefits of GenAI in the workplace, professionals who worked in an organisation that had a GenAI policy generally believed that the policy had supported more comfort in using ChatGPT for work, had improved trust and efficiency, and had provided legal protections. On the other hand, professionals working in organisations without such a policy held mixed views about its value, with approximately half believing it would improve efficiency and provide legal protections, and 40% believing it

would improve trust in the organisation. On these same points, early adopters were generally more optimistic. (Ibid.)

When it comes to policy development around LLMs, organisations appear slow. The same study found that just one quarter (27%) of participants were aware of policies on ChatGPT use at their organisations. The study's authors recognize the importance of engaging early adopters in the process: "Their use of the technologies may allow them to recognize the early benefits of using AI tools, while also allowing them to explain some potential drawbacks and ethical challenges that early use of these technologies has presented (Ibid.)."

2.3 Risks of LLMs for Business

There are several risks associated with the adoption of LLMs for organisations, with the literature commonly mentioning job displacement due to automation, new cybersecurity risks and concerns around data privacy, as well as reliability issues such as algorithmic bias, accuracy, and explainability and the risk of non-compliance with AI regulation.

2.3.1 Identifying Risks of GenAI Development

In a narrative and critical literature review by Wach et al. (2023), researchers conducted an extensive search across academic literature, professional press, and Internet portals on the negative aspects of GenAI development within a management and economics context. They identified various controversies, threats, defects, and disadvantages of GenAI, particularly of ChatGPT. Grouped into clusters, seven main risks were submitted:

1. No regulation of the AI market and urgent need for regulation
2. Poor quality, lack of quality control, disinformation, deepfake content, algorithmic bias
3. Automation-spurred job losses
4. Personal data violation, social surveillance, and privacy violation
5. Social manipulation, weakening ethics, and goodwill
6. Widening socio-economic inequalities
7. AI technostress

Recommendations for managing these risks include regulating the GenAI market, education/retraining of workers in the changing job market, developing systems with user privacy and security in mind, and implementing responsible AI practices and ethical guidelines.

2.3.2 Job Displacement due to Automation

Amidst discussions about LLMs' impressive capabilities and potential productivity gains, there is a significant risk of job displacement. In a paper investigating the potential exposure of the Australian workforce to GenAI (Walkowiak & MacDonald, 2023), researchers quantified, mapped, and analyzed workers' exposure to GenAI and its risks, by measuring their likelihood to manifest within tasks. Their results showed a widespread and massive exposure. It was found that 39% of tasks within jobs are directly exposed to LLMs, accounting for 37% of the time workers spend completing different tasks (Ibid.). Secondly, it was found that 80% of the Australian workforce have 20% of their time allocated to tasks directly exposed to LLMs (Ibid.). The study's risk exposure mapping furthermore showed that LLMs directly expose 12.4% of tasks to privacy risks, 13.7% to cybersecurity risks, 13.6% to breach in professional standards risks, 14.1% to unethical or harmful bias risks, 10.6% to misinformation and manipulation risks, 26% to liability and accountability risks and 9.8% to intellectual property risks (Ibid.). It is important to note, however, that this study covers the whole range of jobs in the Australian workforce, rather than focusing on knowledge work.

Another study on the tasks performed by occupation in the U.S. revealed that around 80% of the American workforce could have at least 10% of their work tasks affected by the introduction of LLMs, while approximately 19% of workers may see at least 50% of their tasks impacted. The projected effects span all

wage levels, with higher-income jobs potentially facing greater exposure to LLM capabilities and LLM-powered software. Analysis furthermore suggested that, with access to an LLM, about 15% of all worker tasks could be completed significantly faster at the same level of quality. This share increases to 47-56% when incorporating software and tooling built on top of LLMs. This implies that LLM-powered software will have a substantial effect on scaling the economic impacts of LLMs. (Eloundou et al., 2023)

The Swedish fintech Klarna is a recent example of an organisation replacing its employees with LLM. At the start of 2024, the bank implemented an LLM-powered AI assistant designed to enhance the shopping and payments experience for 150 million consumers worldwide. The assistant can manage a range of tasks from multilingual customer service to managing refunds and returns. According to the bank, the assistant handled a workload equivalent to 700 full-time staff members and 2.3 million conversations representing two-thirds of the company’s customer service chats in its first month. Moreover, the assistant is available 24/7 in more than 35 languages, is quicker and more accurate than human agents in errand resolution, and scores on par with them when it comes to customer satisfaction. The assistant is estimated to drive a \$40 million USD in profit improvement to Klarna in 2024. (Klarna, 2024)

As LLMs continue to develop, the potential for job displacement is expected to grow. Whereas previously, it was low-skilled and repetitive tasks that were most susceptible to automation by AI (Jetha et al., 2021), developments in LLM capabilities for producing human-like work are increasingly bringing into question the job security of knowledge workers. For instance, lawyers might be surprised to learn that GPT-4 scored in the 90th percentile on the Uniform Bar Exam in July 2022 (Weiss, 2023). The scale and variety of risk posed by LLMs on tasks related to knowledge work should be a consideration for organisations before the widescale implementation of the technology.

2.3.3 Cybersecurity Risks

When considering the adoption of a new digital technology, it is important to also consider the cybersecurity risks associated with it. LLMs are no different in that they too present new vulnerabilities and enable novel ways for cyber attackers to target organizations. To identify the most important cyber risks for organizations, a resource from the Open Worldwide Application Security Project (OWASP) was consulted: OWASP Top 10 Risks for LLM Applications.

OWASP is a nonprofit foundation that works to improve the security of software by enabling organizations to conceive, develop, acquire, operate, and maintain applications that can be trusted. All OWASP projects, tools, documents, forums, and chapters are free and open to anyone interested in improving application security (OWASP, 2024). The OWASP Top 10 for LLM Applications project provides a list of the top 10 most critical risks often found in LLM applications (listed in Table 1). The project highlights their potential impact, ease of exploitation, and prevalence in real-world applications and aims to educate developers, designers, architects, managers, and organizations about the potential security risks when deploying and managing LLMs (Wilson & Dawson, 2023). Examples include new attack vectors specific to LLMs such as prompt injection (LLM01) to gain unauthorized access to confidential data and training data poisoning (LLM03) to impair models with tampered training data, as well as novel ways to perform known cyber-attacks like a denial-of-service attack (LLM04) via LLM.

Table 1: OWASP Top 10 risks and descriptions (Wilson & Dawson, 2023)

OWASP Top 10 Risk	Risk Description
LLM01: Prompt Injection	Manipulating LLMs via crafted inputs can lead to unauthorized access, data breaches, and compromised decision-making.
LLM02: Insecure Output Handling	Neglecting to validate LLM outputs may lead to downstream security exploits, including code execution that compromises systems and exposes data.
LLM03: Training Data	Tampered training data can impair LLM models leading to responses that may

Poisoning	compromise security, accuracy, or ethical behavior.
LLM04: Model Denial of Service	Overloading LLMs with resource-heavy operations can cause service disruptions and increased costs.
LLM05: Supply Chain Vulnerabilities	Depending upon compromised components, services or datasets undermine system integrity, causing data breaches and system failures.
LLM06: Sensitive Information Disclosure	Failure to protect against disclosure of sensitive information in LLM outputs can result in legal consequences or a loss of competitive advantage.
LLM07: Insecure Plugin Design	LLM plugins processing untrusted inputs and having insufficient access control risk severe exploits like remote code execution.
LLM08: Excessive Agency	Granting LLMs unchecked autonomy to take action can lead to unintended consequences, jeopardizing reliability, privacy, and trust.
LLM09: Overreliance	Failing to critically assess LLM outputs can lead to compromised decision making, security vulnerabilities, and legal liabilities.
LLM10: Model Theft	Unauthorized access to proprietary large language models risks theft, competitive advantage, and dissemination of sensitive information.

2.3.4 Ethical Risks and Blind Spots

The increased capabilities of GenAI raise important ethical questions around the technology and its use. Several known ethical issues with AI are also applicable to GenAI, including privacy concerns, a lack of explainability, and algorithm bias. Further ethical risks exist around questions of authorship, authenticity, and plagiarism (Zohny et al., 2023). These risks are especially relevant to the process of developing new LLM policies, whether internally at firms or for new laws for society at large and are discussed further in section 2.3.5 Non-Compliance with AI Regulation.

Firstly, there exist privacy concerns about how AI accesses, uses, and stores the personal data of individuals, and about whether it ensures privacy while handling such sensitive information. LLMs also have significant blind-spots in terms of their susceptibility to catastrophic failure, unreliability (i.e., false or made-up information), and the occasional inability to make elementary logical inferences or do simple mathematics (Floridi, 2023). This behaviour can sometimes be difficult to explain, and this lack of explainability refers to how AI systems often operate like ‘black boxes,’ meaning it is difficult to understand how they arrive at certain decisions or predictions. This lack of transparency can become an ethical issue especially in scenarios where AI decisions significantly impact human lives. AI systems can perpetuate or even amplify biases present in the data they are trained on. This could lead to unfair or discriminatory outcomes, for instance, in hiring decisions, loan approvals, or facial recognition technologies.

Moreover, the automated and effective influence of LLMs at scale poses potential ethical risks, including the risk of user manipulation. Such AI systems have the capability to learn to exploit biases and vulnerabilities in users (Weidinger et al., 2022). While the capability of manipulation can be used beneficially, such as for nudging users to do good, it can also be exploited for manipulative purposes. For example, an LLM-based chatbot may employ manipulative tactics such as peer pressure, emotional guilt trips, and deception to get a user to perform certain actions, or to manipulate consumers' behavior in favor of certain products or services (Klenk, 2023). This manipulation can even happen with users being fully personally autonomous (Klenk & Hancock, 2019).

2.3.5 Non-Compliance with AI Regulatory Requirements

AI regulation is another risk associated with the adoption of LLMs as it could spell heavy fines for organisations found to be non-compliant. The regulation of LLMs is a wicked problem which is highly debated at the international level, and different countries are adopting different approaches. The OECD set out principles in 2019 for trustworthy AI which they have since adapted towards LLMs. The EU and Canada are taking a cross-

sector approach to regulation while the US and UK are employing a more sector-specific approach. (Walkowiak & MacDonald, 2023)

Wicked problems are understood to be loosely formulated and persistent problems subject to redefinition and resolution in different ways over time:

Wicked problems are not objectively given but their formulation already depends on the viewpoint of those presenting them. There is no ultimate test of the validity of a solution to a wicked problem. The testing of solutions takes place in some practical context, and the solutions are not easily undone. (Coyne, 2005)

In the case of LLMs, the technology is highly complex and rapidly evolving. Understanding the nuances of how these models operate, their potential capabilities, and their societal impacts requires expertise across multiple domains, including computer science, ethics, law, and sociology. Moreover, the technology transcends national borders which makes it difficult to coordinate regulatory efforts and enforcement mechanisms across different jurisdictions.

In April 2021, the European Commission proposed an EU regulatory framework for AI. The framework analyses and classifies AI systems according to the risk they pose to users, with different risk levels determining the amount of regulation. Per a June 2023 EU parliament briefing, the categories of risk include unacceptable risk, high risk, limited risk, and minimal risk (Madiega, 2023). The act prohibits certain uses of AI, such as systems that manipulate human behavior in a manner that could potentially cause harm. Specific to GenAI, these systems will have to comply with such transparency requirements as (*EU AI Act, 2023*):

- Disclosing that the content was generated by AI.
- Designing the model to prevent it from generating illegal content.
- Publishing summaries of copyrighted data used for training.

The Act proposes significant fines for non-compliance and is aimed at setting the global standard for AI regulation. Depending on the nature of the infringement, fines can go up to 30 million euros or, in the case of companies, up to 6% of their total worldwide annual turnover from the preceding financial year; whichever is largest. This can apply to instances such as: supplying AI systems considered an unacceptable risk; breaching obligations related to high-risk AI systems; failing to comply with national authorities; and providing incorrect, incomplete, or misleading information to national authorities. (Madiega, 2023)

Given the arrival of new legislation on GenAI, like the EU AI Act, and the associated heavy fines for non-compliance, organisations should be considerate of this risk when adopting LLMs for their business. This places extra importance for organisations to properly understand how the technology is being implemented within their IT infrastructure and how their model functions, to ensure that the necessary guardrails can be in place for compliance.

3 Research Methods

3.1 Research Gap

With the growing prevalence of LLMs, many organisations are gearing up to invest heavily into the technology. Though the technology is a recent advent, professional usage amongst knowledge workers is already widespread and many of them believe that LLM use for work will make them more efficient, help them generate ideas, and improve the quality of their work. Yet, the technology comes with numerous potential risks, including job displacement, threats to data privacy, unreliability, cybersecurity risks, and non-compliance with new AI regulation. The decision to adopt LLMs for knowledge work is thus a significant one as organisations must carefully weigh up the benefits and drawbacks that these models present.

Although many of the benefits and risks associated with LLMs are discussed in the literature, there is still a lack of knowledge on employees' perceptions of LLMs when it comes to enterprise adoption and how they may influence the technology's adoption plans. Specifically, it is not yet clear what the most common use cases of LLMs are for employees, what employees and organisations perceive as benefits and risks regarding LLM adoption, what the current levels of LLM adoption are and organisations' plans for future adoption, and how these align with employee expectations of LLMs. Addressing these questions can shed light on how such perceptions influence organisations' adoption plans, which in turn can inform them on how to develop more effective adoption plans in the future.

3.2 Research Objective

This study seeks to investigate what employees view as the most significant benefits and risks to LLM adoption for their organisations, exploring how and which ones influence organisational adoption of the technology. Additionally, organisations' current LLM adoption and usage policies will be investigated as well as their future adoption plans. The findings of this research should provide insights into strategies for future adoption that align with employee needs and expectations of LLMs. As this research project was completed at the TU Delft in partnership with a global Professional Services Firm (thesis internship provider), the findings aim to achieve both academic and professional relevance.

3.3 Research Scope

The target sector for this research project was the Dutch financial sector. This decision was motivated by the type of clients served by the thesis internship provider and had three notable advantages. Firstly, targeting this sector made the research findings more useful to the internship provider and more relevant to their line of work. Second, it increased the probability of finding the right participants for interview by being able to leverage the internship provider's existing network of financial clients. Third, focussing on the financial sector narrowed the research scope to make the data collected more consistent and more easily comparable.

The research also explored both the employee and organisational perspective to compare them and find out whether and where they match. Where mismatches appeared, these were explored to find out how they misaligned and whether there are conflicts of interest which would affect organisations' LLM adoption plans.

3.4 Research Questions

To achieve the stated research objective, the following questions were formulated.

Main Question:

- How do employees' and organisations' perceived benefits and risks of Large Language Models (LLMs) influence financial organisations' LLM adoption plans?

Sub-Questions:

- SQ1. What are the most common use cases of LLMs?
- SQ2. What are employees' perceived benefits and risks associated with LLM adoption?
- SQ3. What are organisations' perceived benefits and risks associated with LLM adoption?
- SQ4. What are organisations' current levels of LLM adoption and their future LLM adoption plans, and how are they aligned with employee expectations of LLMs?

4 Methodology

This study used a qualitative approach with semi-structured interviews to collect qualitative data. Purposive sampling was used for selecting interview participants, and thematic analysis of interview transcripts was used in the data analysis. The findings were then adapted to the Technology Acceptance Model (TAM) to draw conclusions on how organisations should best handle LLM adoption. The research followed ethical guidelines approved by the university’s Human Research Ethics Committee.

4.1 Qualitative Research

Semi-structured interviews are useful for gaining detailed insights into the perceptions around LLMs within organisations. The main objective is to collect data related to the research sub-questions, such as what participants perceive to be the most significant benefits and risks related to LLM adoption for their organisation. The interviews were conducted in a flexible and interactive way, giving participants the opportunity to expand on their answers and independently voice their experiences. To improve the accuracy and reliability of the analysis, the interviews were audio-recorded for transcription purposes. The revealed details about their perceptions of LLMs were used to answer the research questions.

The analysis of qualitative data is crucial as it is used to investigate the various psychological, social, and contextual factors that influence an individual’s perceptions of LLMs. To achieve this goal, thematic analysis was employed. This involved identifying and coding the interview data according to key themes that emerged from the data, and then organising these themes into a coding framework. The coding framework was developed based on the research questions and sub-questions, refined as the analysis progressed. Lastly, the Technology Acceptance Model was used to integrate the research results within a theoretical adoption framework.

4.1.1 Technology Acceptance Model

The Technology Acceptance Model (TAM) is a widely used theoretical framework developed by Fred Davis in the late 1980’s to understand and predict how users accept and use new information technologies. The model suggests that users' acceptance of a technology depends on their perceived usefulness and perceived ease of use. TAM has been widely used in various fields to predict and explain user acceptance of technologies such as computers, mobile devices, software applications, and websites. It has also served as the foundation for other models and frameworks that explore technology adoption and use. (Davis, 1987)




Figure 3: Technology Acceptance Model flowchart. (Demystifying the Technology Acceptance Model, 2024)

Perceived Usefulness (U) refers to the user's subjective perception of how adopting a particular technology would enhance their job performance or make their tasks easier. If users believe that a technology will be beneficial and improve their effectiveness, they are more likely to accept and use it. Perceived Ease of Use (E) refers to the user's perception of how easy it is to use the technology. It includes factors such as the simplicity of the interface, the ease of learning, and the clarity of instructions. If users perceive a technology as easy to use, they are more likely to adopt it. Both perceived usefulness and perceived ease of use influence users' attitudes toward a technology, which in turn influence their intention to use it. TAM posits that intention to use is the primary determinant of actual usage behavior. External factors like social influence, training, and support can additionally influence users' perceptions and intentions (Davis, 1987).

4.2 Participants

Purposive sampling was used to select interviewees. This sampling approach moves away from any random form of sampling and is used to select respondents that are most likely to yield appropriate and useful information (Campbell et al., 2020). The reason for adopting a purposive sampling strategy is based on the assumption that specific kinds of people may hold different and important views about the ideas and issues at question and therefore need to be included in the sample (Campbell et al., 2020). Participant selection was based on different factors such as the participant working at a Dutch financial institution, having a role proximal to the decision-making process of LLM adoption, having a higher level of education (Bachelor, Master, or PhD), and having work experience relevant to the management of technology. These specifications enabled the study to gain a more consistent understanding of the perceptions surrounding LLM adoption.

Given the novelty of LLMs, it was unfortunately not possible to target one specific role responsible for LLM adoption within each organisation. The objective was therefore to target those roles identified as most involved in the adoption decision-making process. These include, but are not limited to, officers, managers, and advisors in privacy, legal, and risk, as well as heads of Data and AI. It is assumed that the higher the employee's managerial status, the more influence they wield in the LLM adoption decision process. Targeting this wider range of roles offered two advantages. One was an increased probability of finding experts who were available for interview, and the other was the possibility of studying a wider range of knowledge worker types which would make the results more generalizable to the greater business landscape.

Three financial services organisation types were studied: Banking, Insurance, and Pension Fund. These were chosen based on availability as they were the most common organisation types within the internship provider's client network. Moreover, studying multiple organisation types made it possible to compare adoption approaches across them, providing a more representative view of the sector without studying every single one. Overall, these organisations are of interest to the thesis internship provider as part of the tripartite research project agreement, as it is core to the firm's business model of selling professional services to financial services clients.

Table 2 shows an overview of the participants. In total, 18 people were interviewed, sourced via colleagues' professional networks, LinkedIn, in-person contact at a Privacy and GenAI event, and the professional networks of previously interviewed participants. The interviewees' areas of work include Legal, Security, Privacy, Risk, Data, Ethics, AI, and Analytics. Despite inconsistency in the manner of self-reporting level of experience (e.g., some describe years in role, others describe two roles, etc.), a rough average was calculated at 10 years of experience. All participants followed higher education: 3 held a Bachelor degree, 9 held a Master degree (of which, 4 held two Master degrees), and 2 had a PhD.

Table 2: Interview participant overview including organisation, role, experience, and highest education.

Participant Code	Organization Code	Q1. Role	Q2. Years of Experience	Q3. Education
P01	Insurance A	Sr. Legal Counsel AI & Privacy	20 years in Data Processing; 2-3 years in AI Advising	Master (2)
P02	Bank A	Head of Security Analytics	15 years in Analytics	PhD
P03	Bank B	Chief Privacy Officer	9 years in role	Master (2)
P04	Pension Fund A	Data Privacy Officer	5 years in Data Privacy	Master
P05	Pension Fund B	Risk Manager	15 years in role	Master
P06	Pension Fund B	Sr. Risk Manager & Data Protection Officer	6 years as DPO; 21 years as Risk Manager	Master
P07	Insurance B	Data Protection Officer	4 years in role; 23 years in IT & Privacy	Master (2)
P08	Pension Fund A	Privacy Officer	6 years in role; Legal counsel and privacy officer prior	Bachelor
P09	Pension Fund A	Entreprise Risk Manager	2 years in role; 6 years in IT/Data Risk Management	Master
P10	Bank C	Data Privacy Specialist	1,5 years in role	Bachelor
P11	Insurance B	Ethical Advisor	6 years in role	Master
P12	Pension Fund A	IT Risk Manager	5 years in role; 12 years in IT Audit	Master (2)
P13	Insurance A	IT Risk Manager/Business Security Officer	7 years in IT Risk Management (2 years in AI working group)	Bachelor
P14	Bank C	Data Privacy Specialist	1-2 years in role	Master
P15	Insurance B	Entreprise Risk Manager	4 years in Risk Management; 10 years in IT, digitization & AI	Master
P16	Bank A	Head of Data	15 years in Data; 5 in Security & IT Risk	Master
P17	Bank D	Head of Data & AI	8 years in Data & AI	PhD
P18	Bank E	Chief Product Owner AI & Analytics	3 years in AI & Analytics (6 months in role)	Master

The name of each financial organisation was substituted with a code denoting the organisation type, plus an alphabetical identifier (e.g., 'Bank A'). Since participant distribution varied per firm, the identifier was useful during data analysis to indicate which participants worked at the same organisation. In this way, multiples of the same answers could be accounted for so as not to inflate the answer weight of one organisation over another. An overview of the organisation count is given in Table 3. In total, 5 participants worked for an insurance firm, 7 for a bank, and 6 for a pension fund; 9 unique organisations were studied.

Table 3: Organisation count per number of participants and unique organisations.

Organisation Type	No. of Participants	Unique Organisations
Insurance	5	2
Bank	7	5
Pension Fund	6	2
Total:	18	9

4.3 Data Collection & Processing

A total of 18 half-hour interviews were completed. All interviews were held online on the TU Delft Teams environment, followed the same interview structure, and were recorded for transcription purposes. The transcripts were processed, codified, and grouped using the qualitative analysis software Atlas.ti.

4.3.1 Interview Structure

Before starting each interview, an informed consent opening statement was read out and consent for recording the interview was received. The interview format followed a list of 15 questions. When applicable, extra questions were improvised along the way to further explore topics of interest or resolve unclarity within participant answers.

Introduction

Q1. What is your role/position?

Q2. How many years of experience do you have in this field/role?

Q3. What is your highest level of education? (Bachelor/Master/PhD)

LLM Adoption

Q4. Do you personally use LLMs in your work?

If yes, Q4a. Which ones?

If yes, Q4b. How do you use LLMs in your work?

- Q5. Does your organisation currently allow LLMs to be used at work?
If yes, Q5a. Are there any usage policies in place?
If no, Q5b. Why not?
- Q6. Does your organisation have plans for adopting LLMs in the future?
If yes, Q6a. What is the timeline on these plans?

Employee & Organisational Perspectives

- Q7. What do you perceive as the biggest benefits to yourself when using LLMs for your own work?
Q8. What do you perceive as the biggest risks to yourself when using LLMs for your own work?
Q9. What do you perceive as the biggest benefits to your organisation when using LLMs in your work?
Q10. What do you perceive as the biggest risks to your organisation when using LLMs in your work?

Perspective Alignment

- Q11. Is there a risk/benefit that could be beneficial to the user but pose a risk to the organisation?
Q12. Is there a risk/benefit that could be beneficial to the organisation but pose a risk to the user?

Most Excited/Afraid About

- Q13. What are you most excited about with regards to LLM adoption within your organisation?
Q14. What are you most worried/afraid of with regards to LLM adoption within your organisation?

User Manipulation

- Q15. Is there a risk of manipulation of the user when using LLMs in your work?

Introduction questions (Q1-Q3) gathered participant background information. LLM adoption questions (Q4-Q6) sought to understand if LLMs are allowed by the organisation, the current level of adoption, LLM use cases, any usage policy, and future adoption plans including timelines. Perspective questions (Q7-Q12) explored the perceived benefits and risks for both employees and organisations to reveal which ones are the most common, and whether these are aligned to identify potential conflicts of interest affecting the adoption process. Participants were also asked to describe what they are most excited and afraid about regarding LLM adoption within their organisation (Q13-Q14) to infer their wishes/expectations of LLMs. Finally, participants were asked whether they perceived a risk of user manipulation through their use of LLMs (Q15). This question is a remnant of the study's original research direction and was therefore not directly relevant to answering the final set of research questions.

After each interview was completed, participants were thanked for their time and were promised a one-page summary of the results to be sent by email after analysis as a token of appreciation.

4.3.2 Interview Transcription

Transcripts were automatically generated and downloaded from the Teams environment. Each interview was replayed to rectify mismatches between recording and transcript, and to remove any recorded stutters, repeated words, filler words, and information not relevant to the research such as greetings or small talk before or after the interview. Any sensitive or personal information such as the names of participants, their organisation, and competitors were removed or replaced with an anonymous label (e.g., '[bank]').

4.3.3 Transcript Coding & Grouping

The anonymized transcripts were uploaded and codified in Atlas.ti, a qualitative analysis software tool. When needed, comments were created as personal reminders to clarify ambiguities in participant responses or to explain important discoveries. All throughout the coding process, regular reviews of the codes were performed in the code manager, removing any redundant codes, and merging similar codes together. Merge examples include combining 'Efficiency' with 'Productivity' to form 'Efficiency/Productivity', or 'Improve Writing' with 'Improve Grammar' and 'Improve Translation' to form 'Improve Writing/Grammar/Translation'. In short, any

code containing a slash or multiple synonymous descriptors in its name is a combined code. Repeated codes were deleted.

Grouping served to facilitate the search and frequency analysis within Atlas.ti. For instance, by filtering all codes by the group 'Risks', codes describing perceived risks are listed and can be ordered by frequency. Each code was assigned to one of the following groups: 'Role', 'Adoption', 'Usage Policy', 'Benefits', 'Risks', 'Descriptive', and 'Misalignment of Perspectives'. Most of the groups relate to an interview question category. 'Role' contains the introduction questions (Q1-Q3) plus the code "Role Description", which codifies extra information provided in the interview on role responsibilities. Similarly, 'Adoption' includes all LLM adoption questions (Q4-Q6) related to current and future LLM adoption, except for Q5a which has its own group 'Usage Policies' as this was an area of particular interest that was significantly different to warrant a separate group. On the other hand, 'Benefits' and 'Risks' capture overall participant perceptions by grouping all benefits and risks codes mentioned across interview questions. This same scope applies to 'Descriptive', however, this group contains neutral codes (i.e., neither a benefit nor a risk) which describe the responses, such as "State-of-the-Art" or "Financial Data". Although not directly related to a specific interview question, this group served to create an overview of the various themes found across the interviews which could quickly be reviewed from the code manager during analysis. Finally, the group 'Misalignment of Perspectives' includes Q11 and Q12 on perspective alignment between employee and organisation, as well as codes specific to this topic such as "Change Management" and "AI Aversion".

4.4 Ethical Considerations

During the study, participant data was collected such as role, years of experience, and education level. In this light, the research project has followed ethical guidelines and obtained approval from the Human Research Ethics Committee (HREC) of Delft University of Technology. Participants were provided with information regarding the study, their rights as participants, and any potential risks or harms. Informed consent was obtained ahead of each interview and the option to withdraw from the study at any point without reason was explicitly stated. Participant codes (e.g. 'P01') ensure participant anonymity, and personal or firm sensitive information was removed from the data before analysis to protect the participants' identities. The data collected was stored in the university's secure environment per the HREC's recommendations to maintain security and confidentiality.

5 Analysis of LLM Adoption Perceptions

5.1 Benefits Overview

Across all interview transcripts, 42 benefits were mentioned a total of 245 times and thematized into eight themes: 'Assistance' for codes that assist an employee in their work tasks (e.g., "Information Search"), 'Business' for codes that are directly beneficial to the organisation's bottom line (e.g., "Cost Savings"), 'Creative' for codes that have a dimension of creativity (e.g., "Brainstorming/Inspiration"), 'Efficiency' for codes that improve the efficiency of tasks/processes (e.g., "Automate Repetitive Tasks"), 'Higher Quality' for codes that describe improvements to quality within work outputs (e.g., "More Precise/Accurate"), 'Literary' for codes that have a literature component (e.g., "Text Generation"), 'Sentiment' for codes related to employee sentiment about LLMs (e.g., "Amazement"), and 'Other' for two codes that did not fit any of the previous themes, namely "Increased Understanding of LLMs" and "More Human-to-Human Work". These themes were chosen based on the commonalities between the benefits codes. The benefit codes and related themes are found in Table 4, ordered from most to least frequently mentioned.

Table 4: Thematized benefit codes ordered by frequency, including percentage share of total benefit codes.

Benefit	Theme	Frequency	Share
Efficiency/Productivity	Efficiency	40	16.3%
Automate Repetitive Tasks	Efficiency	18	7.3%
Improve Writing/Grammar/Translation	Literary	13	5.3%
Information Search	Assistance	13	5.3%
Customer Service Improvement	Business	12	4.9%
More Focused/Interesting/Fun Work	Creative	9	3.7%
Cost Savings	Business	8	3.3%
Maintain/Improve Quality of Work	Higher Quality	8	3.3%
Summarise Documents	Literary	8	3.3%
Summarise Meetings/Calls	Efficiency	8	3.3%
Improve Programming	Efficiency	7	2.9%
More Precise/Accurate	Higher Quality	7	2.9%
Text Generation	Literary	7	2.9%
Creativity Benefit	Creative	6	2.4%
Email Composition	Literary	6	2.4%
New Opportunities	Business	6	2.4%
AI/Personal Assistant	Assistance	5	2.0%
Brainstorming/Inspiration	Creative	5	2.0%
Creating Presentations	Creative	5	2.0%
Good Starting Point	Assistance	5	2.0%
Structuring Documents	Literary	5	2.0%
Amazement	Sentiment	4	1.6%
Image/Video Generation	Creative	4	1.6%
Optimism wrt LLMs	Sentiment	4	1.6%
Competitive Advantage	Business	3	1.2%
Conversational Analysis	Literary	3	1.2%
Learning Opportunity	Creative	3	1.2%
AI Democratization	Business	2	0.8%
Cultural Benefit	Sentiment	2	0.8%

Data Quality Improvement	Higher Quality	2	0.8%
Experimenting	Creative	2	0.8%
Getting a Concrete Answer	Assistance	2	0.8%
Increased ROI (Return on Investment)	Business	2	0.8%
Optimization	Business	2	0.8%
Provides Employee Independence	Creative	2	0.8%
Consistency	Higher Quality	1	0.4%
Increased Understanding of LLMs	Other	1	0.4%
More Human-to-Human Work	Other	1	0.4%
More Personal Human-Computer Interaction (HCI)	Assistance	1	0.4%
Perform Analysis	Assistance	1	0.4%
Scalability	Business	1	0.4%
Versatility	Assistance	1	0.4%

The five overall most common perceived benefits for the adoption of LLMs are efficiency/productivity gains (16%), the automation of repetitive tasks (7%), facilitating information search (5%), improving writing/grammar and translation (5%), and customer service improvement (5%). Some of the benefits only mentioned once (0.4%) include scalability, consistency, versatility, and an increased understanding of LLMs.

Customer service improvement is the most common perceived business benefit of LLMs. Participants mentioned that LLMs can automate customer support to decrease waiting time, perform conversational analysis to predict customer questions and optimize customer service. Moreover, LLMs help bring customer information closer to the hands of customer service representatives by facilitating information search, and P18 describes how LLMs help create a “hyperpersonal” AI assistant capable of providing the right offerings at the right time to their commercial clients.

Table 5: Frequency analysis of ‘benefits’ themes, ordered by total frequency of codes per theme.

Benefit Themes	No. of Unique Codes	Total Theme Frequency	Share
Efficiency	4	73	29.8%
Literary	6	42	17.1%
Creative	8	36	14.7%
Business	8	36	14.7%
Assistance	7	28	11.4%
Higher Quality	4	18	7.3%
Sentiment	3	10	4.1%
Other	2	2	0.8%

As shown in Table 5, ‘Efficiency’ (30%), ‘Literary’ (17%), and ‘Creative’ (15%) are the largest themes by number and share of total mentions. Meanwhile, the themes ‘Creative’ and ‘Business’ are largest in terms of their number of unique codes (8 codes).

5.2 Risks Overview

Across all interview transcripts, 50 risks were mentioned a total of 296 times and thematized into six themes: ‘AI Risk’ for risk codes that are specific to or have come about by AI systems (e.g., “Overreliance on LLM”), ‘Business’ for codes that are directly detrimental to the organisation’s bottom line (e.g., “Reputational Risk”), ‘Cybersecurity’ for cyber risk codes (e.g., “Data Leakage/Breach”), ‘Human/Social’ for codes that impact employees or the broader society (e.g., “Job Loss due to Automation”), and ‘Operational’ for codes that are detrimental to business operations (e.g., “Unreliability”). These themes were chosen based on the

commonalities between the risk codes. The risk codes and related themes are found in Table 6, ordered from most to least frequently mentioned.

Table 6: Thematized risk codes ordered by most frequent, including percentage share of total risk codes.

Risk	Theme	Frequency	Share
Lack of Quality Control/Output Validation	AI Risk	29	9.8%
Inputting/Exposing Sensitive Data	Cybersecurity	27	9.1%
Data Leakage/Breach	Cybersecurity	20	6.8%
Job Loss due to Automation	Human/Social	18	6.1%
User Manipulation*	AI Risk	17	5.7%
Bias	Human/Social	14	4.7%
Overreliance on LLM	AI Risk	13	4.4%
Unreliability	Operational	12	4.1%
Explainability/Black Box	Human/Social	11	3.7%
Regulatory Compliance	Business	7	2.4%
Distinguishing Real from Fake	AI Risk	6	2.0%
Loss of Control	AI Risk	6	2.0%
Negative Impact on Customers	Business	6	2.0%
Reputational Risk	Business	6	2.0%
Hallucination/Delusion	AI Risk	5	1.7%
IP Infringement	Business	5	1.7%
Overkill/Using Wrong Solution	Operational	5	1.7%
Shadow Security/Workarounds	Cybersecurity	5	1.7%
Complacency	AI Risk	4	1.4%
Creativity Risk	AI Risk	4	1.4%
Ethical Risk	Human/Social	4	1.4%
Lack of Transparency	Business	4	1.4%
Loss of Expertise/Tacit Knowledge	Business	4	1.4%
Phishing/Spam	Cybersecurity	4	1.4%
Programming Risk	Operational	4	1.4%
Unknown Risks	AI Risk	4	1.4%
Work Becomes Dull/Uninspiring	Human/Social	4	1.4%
Data Security	Cybersecurity	3	1.0%
Environmental Sustainability	Human/Social	3	1.0%
Lack of Critical Thinking/Reflection	AI Risk	3	1.0%
New Cyber Threats	Cybersecurity	3	1.0%
New Risk Profile	AI Risk	3	1.0%
Poor Data Quality	Operational	3	1.0%
Skill Deterioration	Human/Social	3	1.0%
Widening Competitive Divide	Business	3	1.0%
AI Aversion	AI Risk	2	0.7%
Bankruptcy	Business	2	0.7%
Cultural Risk	Human/Social	2	0.7%
Fraud	Cybersecurity	2	0.7%
Impersonal/Loss of Personal Touch	Human/Social	2	0.7%
Lack of Accountability	Human/Social	2	0.7%

Lawsuit	Business	2	0.7%
Ransom	Cybersecurity	2	0.7%
Sabotage	Cybersecurity	2	0.7%
Improper Implementation	Operational	1	0.3%
Lack of Safety Measures	Cybersecurity	1	0.3%
Reduced Job Satisfaction	Human/Social	1	0.3%
Single Point-of-Failure	Cybersecurity	1	0.3%
Social Engineering	Cybersecurity	1	0.3%
Training Data Poisoning	AI Risk	1	0.3%

The five overall most common perceived risks for the adoption of LLMs are the lack of quality control/output validation (10%), inputting or exposing sensitive data (9%), the risk of a data leakage/breach (7%), job loss due to LLM automation (6%), and user manipulation (6%). It should be noted that the number of times user manipulation was mentioned as a risk is inflated due to Q15 on the risk of user manipulation. Some of the least commonly mentioned risks include LLM training data poisoning, social engineering, single point-of-failure, and reduced job satisfaction (0.3% each).

The top three most common perceived risks are cybersecurity and AI risks. When relating these to the OWASP Top 10 for LLM applications, it is seen that lack of quality control/output validation is LLM02 'Insecure Output Handling', and inputting/exposing sensitive data together with data leakage/breach is LLM06 'Sensitive Information Disclosure'. It is also interesting to note that training data poisoning (LLM03) was mentioned only once and did not arise more frequently throughout the interviews, given it is a risk specific to LLM technology. Similarly, LLM01 'Prompt Injection', LLM04 'Model Denial of Service', LLM05 'Supply Chain Vulnerabilities', LLM07 'Insecure Plugin Design', and LLM10 'Model Theft' were never mentioned which begs the question about whether these risks are not well known, not a perceived concern, or not relevant.

Table 7: Frequency analysis of 'risks' themes, ordered by total frequency of codes per theme.

Themes (Risks)	No. of Unique Codes	Total Theme Frequency
AI Risk	13	97
Cybersecurity	12	71
Human/Social	11	64
Business	9	39
Operational	5	25

Table 7 summarizes the risk themes in terms of the number of unique codes associated to each theme and the total frequency of the theme. It was found that 'AI Risk' (97 mentions), 'Cybersecurity' (71 mentions), and 'Human/Social' (64 mentions) are the largest themes by both associated code count and in their number of unique codes.

5.3 Current LLM Adoption

This section presents findings on whether organisations currently allow LLMs to be used at work, the reasons for banning LLMs, whether employees use LLMs, and the most common use cases.

5.3.1 Authorization of LLMs

Analysis of responses to Q5 on whether LLMs are authorized by the organisation revealed that 14 out of 18 participants are allowed to use LLMs for their work, with 3 of these 14 authorizations classified as 'limited usage', such as LLMs being tested in pilot and/or not being made available for all employees/departments. The remaining 4 participants stated that their organisation prohibits LLMs in the workplace.

The responses given to Q5b on why LLMs are not allowed explain the motivations behind the managerial decision to prohibit LLM access. The general sentiment from organisations prohibiting LLMs is that the risks associated with its use are too large, specifically regarding the exposure of business/confidential information or personal data from prompts inputted into models. These organisations do not yet feel comfortable to allow their employees to use LLMs and are waiting to investigate the technology further, to improve their understanding of it and implement the necessary safeguards for the protection of employees and customers.

- **P02:** "It is a technology that is still being understood and the management board does not feel like it has the necessary safeguards around LLMs to use that in a safe way for employees and customers."
- **P04:** "It poses some risks with regard to what kind of information might be put there."
- **P05:** "For the time being, with all the uncertainties and the risk appetite of the company, it's been put on hold. Not forever, but we still want to investigate further before we [proceed]."
- **P16:** Participant was not directly asked why LLMs are prohibited, but the reason is assumed to be the same as for P02 since both participants work at Bank A.
- Additionally, **P14** explained that although Copilot had recently been authorized at Bank C, it was previously blocked because management "didn't want people to put all kinds of information in there, especially business/confidential information or personal data."

Seeing as the distribution of participants working at different organisations is not uniform, it is important to understand to what extent LLMs are authorized per firm. Table 8 summarizes whether each organisation represented by the interviewees allowed LLMs at the time of interview. Only 2 out of the 9 organisations prohibited LLMs for use at work: Bank A and Pension Fund B.

Table 8: LLM authorization per organisation (Q5).

Organisation	Q5. LLMs Allowed?
Bank A	No
Bank B	Yes
Bank C	Yes
Bank D	Yes
Bank E	Yes
Insurance A	Yes
Insurance B	Yes
Pension Fund A	Yes
Pension Fund B	No

During the analysis for Pension Fund A, it was noticed that P04 stated LLMs were not allowed at work while three other colleagues said they were. This discrepancy could be explained in a few ways. Firstly, P04 was one of the first participants interviewed and the usage policy could have changed by the time P08, P09, and P12 were interviewed. Second, the question on whether LLMs are allowed could have been interpreted as applicable to all employees or specifically to the participant's role. P08 stated that the firm is working on a couple of pilots which are not available to all employees, thus, it would follow that P04 did not have access to LLMs in their role. Lastly, P04 could have been mistaken in their understanding of whether LLMs were allowed. For these reasons and given that all three of P04's colleagues stated otherwise, it is assumed that Pension Fund A authorizes LLMs for use at work.

All in all, the results seem to indicate that only a minority of financial organisations prohibit employees from using LLMs for their work due to not having a good enough understanding of LLMs, not having the proper

safeguards in place for its use at work, the risk of exposing sensitive data that is inputted in the LLM, and the organisation's lower risk appetite.

5.3.2 Usage Policies

Looking at existing policies gives insight into the current organisational levels of LLM adoption by learning about the details of what is and is not allowed, corporate guidelines and recommendations for LLM use, and other specific rules. This analysis is important for identifying how organisations can change their policies to accommodate the wishes/expectations of LLM usage at work in their future adoption plans.

In the interviews, participants were asked to comment on any LLM usage policies that their organisation had in place. Analysing their responses per organisation led to the summary in Table 9.

Table 9: Summary of LLM usage policies per organisation.

Organisation	Q5a. Usage Policy (summarized)
Bank A	- Not allowed.
Bank B	- Open to specific people for specific purposes as part of testing.
Bank C	- Do not input company/confidential information or personal data.
Bank D	<ul style="list-style-type: none"> - Depending on the use case, certain LLMs may be used as part of a certain process, like as an assistant or for automation. - In other use cases, a "bring your own LLM" policy is followed whereby employees are free to choose the LLM they prefer. - If the LLM is hosted outside the company environment, employees are not allowed to share sensitive data with the LLM.
Bank E	<ul style="list-style-type: none"> - LLMs must be used in a secure environment owned by Bank E. - Publicly available LLMs may only be used if the input data is unrelated to the bank. - Employees have unrestricted access to all LLMs except for ChatGPT. Access to ChatGPT has been completely restricted but will reopen with managed access after the implementation of a new security measure that prevents employees from inputting restricted data. - Officially, there are two sets of LLM usage guidelines: one for programmers building AI solutions and another for general employees. The latter is a knowledge and awareness document including considerations to make before using the tool, what shouldn't be done, and where to go if you have any questions, doubts, or concerns about the usage. Both have been approved by the bank's decision-making organs around privacy, security, etc.
Insurance A	<ul style="list-style-type: none"> - A list of do's and don'ts, such as to be aware the data inputted into the model or always checking the model output. - A specific policy on how to use ChatGPT-based chatbots for end-users. - A formal policy for the use of LLMs for software development.
Insurance B	<ul style="list-style-type: none"> - No specific policy in place for the use of LLMs. It is therefore possible to install and use any tool without restriction. - For general employee use, the firm is still working on some broad guidelines. Currently, an employee wanting to use any form of personal data must do an online assessment with filters for the LLM type and use case being requested. - The firm does have guidelines on what data can be used in prompts, how to inform colleagues about LLM usage in produced work, etc. It not so much a policy as it is customary / a cultural agreement.
Pension Fund A	<ul style="list-style-type: none"> - Not available to all employees (pilot phase) - No policy has not been formalised yet, but the firm is in the final stages of

	formulating its LLM policy.
Pension Fund B	<ul style="list-style-type: none"> - A memo has been written stating that LLMs are currently not allowed, but that the firm is working on making them available. - The organisation does not allow the free use of LLMs in the company environment: an employee must receive special clearance from the Security, Architecture (IT), and Risk & Privacy teams. When a request for use is made, the employee receives the rules and regulations stating which models can be used and in which ways. - The company's system is blocked from the use of ChatGPT. - The usage policy has very specific rules prohibiting the use of any personal/company information in model prompts.

Overall, the policies can be categorised in three ways. The first is a complete ban on LLMs which aims to eliminate the risks of LLM usage, or at least temporarily until new LLM policies are finalised and implemented. The second and most common type is a partial ban, where certain LLMs are prohibited but not others. In some cases, the ban is partial because it applies to certain roles and departments whereas others have received or can request LLM access. Often, this has to do with the testing of an LLM pilot for which the software team has access to the LLM for development purposes. The third category is free use and is the least common type. In these instances, firms trust their employees to use LLMs responsibly and rely on their common sense.

The rules and guidelines in organisations' usage policies all mention that inputting sensitive data is not allowed, and most mention that model outputs must be checked before use. This finding aligns very well with the top perceived risks of lack of quality control/output validation, inputting/exposing sensitive data, and data leakage/breach found previously in 5.2 Risks Overview, showing that policies are trying to address and mitigate this risk. Beyond this, most policies provide guidance on how to use the model and where to go for assistance. Some firms also have separate policies specific to the use of LLMs for software development.

The policies are generally paper policies which provide rules or recommendations on how models are to be used, rather than blocking employees through technical means from inputting certain information into the models. Among others, this is the approach of Bank E which "relies on trust in their employees and does a lot of awareness around responsible use." P18 said that the bank has been doing this for the last couple of years on a number of tools and it is observed that employees are well aware of the risks. P12 also mentioned that "besides the official policy, employees are expected to use their common sense."

Finally, some concerns were raised about usage policies not being well known or well communicated to employees. P05 was doubtful as to whether everyone knew about Pension Fund B's memo on LLM usage and exactly what it allows, while P11 said that the policy is "not broadly known or communicated because it's a moving target."

5.3.3 Employee LLM Usage & Most Common Use Cases

Exactly half of the participants (9/18) said that they use LLMs in their own work while the other half stated that they do not. The motivation for the latter half not to use LLMs is due to a lack of perceived benefits for their work. For example, P08 is a privacy officer who said that there is a lot of private information used in their work, therefore, they do not see the benefit of using LLMs which risk leaking data. Similarly, P11 said that although it is good a reproducing a lot of information, it does not help them in their work as an ethical advisor. Interestingly, P16 uses LLMs for work despite an organisational ban, describing their use as being part of a broader work culture within the engineering team to use workarounds to access LLMs for their programming work.

- **P08:** "I don't really see the benefits at this moment in my work because there's a lot of private information in what I work with. So, I'm not taking the risk."

- **P11:** “It’s good at reproducing a lot of information but it’s not really helping me do my work any further. I haven’t played around with them much because my work involves talking to people, helping them structure their work in certain ways.”

Table 10 summarizes the answers of the 9 LLM users, including which LLMs they use (Q4a.) and what they use it for (Q4b.). Interestingly, most of these users have higher numbered participant codes which indicates that they were in the latter half of the interviewed participants. This observation could already be an indication of a high pace of LLM adoption within the industry, meaning that the time elapsed between earlier and later interviews could be a significant factor in the participants’ responses on their usage of LLM.

Table 10: LLMs used by participants and LLM use cases.

Participant	Q4a. Which LLMs do you use?	Q4b. How do you use LLMs?
P01	Organisation’s private ChatGPT, Dall-E, Bing Content Creator	Email Composition; Creating Presentations; Image/Video Generation.
P09	ChatGPT (used the most), Cloud (like ChatGPT but focused on comparing PDFs), Canva	Brainstorming/Inspiration; Email Composition; Text Generation; Improve Writing/Grammar/Translation; Image/Video Generation; Creating Presentations.
P10	Copilot 365, GitHub Copilot	Structure Documents; Create Presentations; Automate Repetitive Tasks.
P13	Organisation’s private ChatGPT	(Did not ask)
P14	Copilot	Information Search; Text Generation; Brainstorming/Inspiration.
P15	Copilot	Information Search; Summarise Meetings/Calls; Document Structuring.
P16	Google Bart, ChatGPT, and Copilot	(Did not ask)
P17	Gemini, GPT4 (most used), GPT3.5, LLaMa, Mistral, Orca	Improve Writing.
P18	Multiple GPT models including Copilot, Midjourney, Grammarly, and People (a translation service, also for documents, as well as it teaches you how to write better)	Improve Grammar/Writing/Translation; Text Generation; Image Generation; Email Composition; Structuring Documents; Creating Presentations.

A frequency analysis of Q4a responses (Table 11) revealed that the most common LLMs used are Microsoft’s Copilot (6 mentions), OpenAI’s ChatGPT (4 mentions), and private versions of ChatGPT (2 mentions), with the rest of the LLMs only being mentioned once. In total, 14 unique LLMs were recorded with a combined total of 23 mentions. With Microsoft being a multiyear-long investor in OpenAI’s ChatGPT and given that Copilot is based on the same GPT technology as ChatGPT (Spataro, 2023), it is clear that Microsoft is a significant player in the current adoption of LLMs at Dutch financial institutions. When adding Bing Content Creator (another Microsoft LLM application) to this list, Microsoft’s share comes out to 13/23 or 57% of all LLMs used by participants.

Table 11: Participant LLM usage, ordered by most frequently mentioned (Q4a).

LLM Type Used	Frequency	Share
Microsoft Copilot	6	26.1%
OpenAI ChatGPT/GPT	4	17.4%
Private ChatGPT	2	8.7%
Dall-E	1	4.3%

Bing Content Creator	1	4.3%
Cloud	1	4.3%
Canva	1	4.3%
Google Bart	1	4.3%
Google Gemini	1	4.3%
Meta LLaMa	1	4.3%
Mistral	1	4.3%
Midjourney	1	4.3%
Grammarly	1	4.3%
People	1	4.3%

A frequency analysis of Q4b responses (Table 12) revealed that the most common LLM use cases by participants include preparing presentation slides (16%), text generation (12%), email composition (12%), and structuring documents (12%). Observing the participant use cases reveals that the participants are indeed knowledge workers as they must perform relatively complex tasks which are creative, analytical, and literary in nature. These use cases and corresponding themes align well with the definition of knowledge work from 2.1.2 Knowledge Work, which describes a knowledge worker as having their major work tasks involve the creation of new knowledge or the application of existing knowledge in new ways, using information technology, theoretical knowledge, creativity, and analytical skills.

Table 12: Participant LLM use cases, ordered by most frequently mentioned (Q4b).

LLM Use Case	Theme	Frequency	Share
Creating Presentations	Creative	4	16.0%
Text Generation	Literary	3	12.0%
Email Composition	Literary	3	12.0%
Structuring Documents	Literary	3	12.0%
Image/Video Generation	Creative	3	12.0%
Improve Writing/Grammar/Translation	Literary	3	12.0%
Information Search	Assistance	2	8.0%
Brainstorming/Inspiration	Creative	2	8.0%
Automate Repetitive Tasks	Efficiency	1	4.0%
Summarise Meetings/Calls	Efficiency	1	4.0%

The use cases are thematized with the same benefit themes from section 5.1 Benefits Overview. It is observed that the most common themes are 'Literary' and 'Creative'. Although the most common use case is 'Creating Presentations' which is themed as 'Creative', the theme itself only has 9 mentions as compared to the most common theme by total frequency, 'Literary' with 12 mentions (Table 13).

Table 13: Participant LLM use case themes, including number of codes per theme and total frequency.

Themes	No. of Codes	Frequency
Literary	4	12
Creative	3	9
Efficiency	2	2
Assistance	1	2

5.4 Future LLM Adoption

This section explores organisations' future LLM adoption plans. Participant answers make it possible to better understand how organisations are managing this adoption, including motivators for adoption, specific adoption strategies, steps being taken towards it, and the challenges holding back adoption plans.

Participants were first asked Q6 on whether their organisation has plans to adopt LLMs in the future. Of the 18 participants, 13 expressed that their firm did, 3 were unclear or unsure, and 2 said that their firm did not. Viewed per organisation, it was found that 7/9 firms have future adoption plans, one is unclear, and another does not. Many of the responses also included mentions of LLM pilots being run to assess the technology and its usefulness within the organisation. 14/18 participants mentioned that their organisation had one or more pilot running at the time of interview. In terms of organisations, this is equivalent to 7/9 firms running a pilot. Specifically, all firms had a pilot active with two exceptions: Insurance A and Pension Fund B. This shows that most firms have taken a first step towards LLM adoption. These findings are summarised in Table 14.

Table 14: Participant responses on future LLM adoption plans (Q6), including mentions of active LLM pilots.

	Future adoption plans?		Pilot?	
	Participants	Organisations	Participants	Organisations
Yes	13	7	14	7
Unclear	3	1	0	0
No	2	1	4	2

The interview answers to Q6 on future LLM adoption plans are summarised in Table 15 and a deeper analysis reveals the underlying motivations for adopting LLMs or for holding off instead.

Table 15: Summary of future LLM adoption plans per organisation.

Organisation	Q6. Future adoption plans (summarized)
Bank A	Yes. Copilot is in pilot to help programmers write better code. Once in production, it will become available to all employees (not just engineers).
Bank B	Yes. Running a pilot to determine what the bank wants to do with the LLM, and the intention is there to adopt more LLMs in the future. If they are useful enough and the risks can be managed properly, then the drive is there to use LLMs more.
Bank C	Unclear. The bank is currently using Copilot, and possibly more LLMs in the future.
Bank D	The bank has someone focusing on LLMs to make sure that they have the state-of-the-art available. Whenever something better comes along, it gets tested and used.
Bank E	In the future, the bank will have managed access to more LLMs (including ChatGPT).
Insurance A	Yes. The firm has plans to use other, more specific LLMs, and is currently creating solutions based on the technology of ChatGPT for very specific purposes within the company. Within P13's department specifically, not that many models have been adopted yet. They are designing the first use case to make the work of security documentation easier.
Insurance B	Yes. Insurance B will adopt LLM applications from Microsoft. The firm is developing the use of GPT for analysing customer help desk conversations. The firm is also looking to use LLMs to analyse emails and predict customer questions. Another use case is internal knowledge models that can be interacted with to search for documents and ask questions about them. There are about 10 to 12 use cases at different stages of development in different parts of the organisation.
Pension Fund A	Yes. There are a couple of pilots including talk about a Bing pilot (chatbot). The firm is working on allowing LLMs in a controlled environment. Not a private LLM but rather via APIs to access an external LLM.

Pension Fund B	No. Pension Fund B is not that far yet.
----------------	---

5.4.1 Future LLM Adoption Capabilities

Based on the answers to Q6, it is difficult to make a comprehensive account of the future capabilities of the LLMs looking to be adopted by organisations in the future. This is because some organisations are still investigating use cases and therefore the requirements for LLMs is still unclear, while other answers do not provide enough details about what the LLMs are specifically expected to achieve. Nevertheless, some of the mentioned uses are as follows:

- Using Copilot to help programmers write better code.
- Using GPT to analyse conversations between customers and the help desk staff by transcribing data speech-to-text and making summaries, gathering new insights from this data.
- Using GPT to analyse emails and predict customers' questions.
- Using Bing as a general chatbot assistant to answer employees' questions.
- Using GPT on internal knowledge models by asking the LLM a question for it to return the correct documents from the company database.

5.4.2 Adoption Decision Process and Strategies

Firstly, many firms want to explore the possibilities of LLMs through pilots to research their added value. An emphasis is placed on risk assessments to make sure the risks are acceptable before continuing and so that the necessary safeguards can be implemented. Ultimately, it is a strategic decision on whether the organisation believes that the benefits outweigh the risks.

P04 and P16 mention that Bank A's new LLM tools are being held back from going into production until usage policies are in place. P16 furthermore explains that there is pressure from the engineering team to adopt LLMs:

The engineering culture is a driving force for adopting LLMs. Due to the ban, some engineers are creating workarounds to access LLMs which is a security issue. Others are leaving the company for other places where they do have LLMs available to them.

According to P07, competition is a major factor for Insurance B to adopt LLMs: "There's a lot of pressure within the organisation to adopt it as soon as possible to be able to show that we do something with AI too, because one of our major competitors uses it already." The firm seems to be taking a 'Smart Follower' approach to adoption: "I saw a position paper on our AI wherein we call ourselves a 'smart follower'. We're not ahead in the development, but we are willing to adopt it as soon as possible."

Pension funds seem to be more risk averse when it comes to LLM adoption. Pension Fund A is "under a lot of scrutiny from the regulators" due to being in the pension asset management business, according to P09. Similarly, P05 explains that Pension Fund B is a more risk-aware company by virtue of being a pension fund: "We're investing long-term, so you don't want to risk it all because of some application that you don't know anything about yet." Nevertheless, competition remains a driver for adoption because "if you block it completely, you will miss out as well, because all pension funds will start using it in the end. You would be stupid to lose the opportunity" (P05).

As for the decision to develop an LLM in-house or to buy it from a vendor, P18 explains that Bank E does not just consider whether it is technically possible for their development team to build the LLM, but also the cost of ownership, the effort required to keep it up-to-date, and how many products the team can manage at once. Use cases for the LLM are identified in terms of impact to the business and effort to produce, always weighing

both factors. To manage this complexity, multiple decisions are made throughout the process on whether to proceed with development.

Overall, most organisations have already implemented pilots to assess the risks of LLMs prior to adoption. There is pressure on organisations to adopt quickly due to competition, with Insurance A taking a ‘Smart Follower’ approach to reduce adoption risk while keeping up with competitors. Meanwhile, Bank A is under pressure to adopt as the lack of access to LLMs for programmers is creating a culture issue. Programmers are resorting to using workarounds to access these tools for their work and some are leaving the company for positions at competing firms that do have LLMs available. Barriers to future adoption plans include the lack of usage policies in place, and especially for pension funds, regulatory scrutiny, and a lower risk appetite due to the long-term nature of their investments. Finally, the decision to license or develop LLMs is an iterative process for Bank E which considers the costs of ownership and maintenance, as well as the impact on business vs the effort to produce.

5.5 Employee Perspective

This section analyzes the interview responses on employees’ perceived benefits and risks associated with LLM adoption by way of a frequency analysis of the risk codes associated with Q7 and Q8.

5.5.1 Employee Perceived Benefits

To understand the employee perspective on LLM adoption benefits, the responses to Q7 were analysed and summarised in Table 16.

Q7. What do you perceive as the biggest benefits to yourself when using LLMs for your own work?

The five most common perceived employee benefits when it comes to using LLMs for work are increased efficiency/productivity (11%), improved writing/grammar/translation (9%), easier information search (9%), assistance with structuring documents (9%), and automating repetitive tasks (7%). Except for document structuring, these most common perceived employee benefits were also found in the top five overall most common perceived benefits. Interestingly, the answer ‘(No Benefit)’ is due to P06 not perceiving any employee benefits of using LLMs in their work, explaining that they already have easy access to the information needed for their work in the risk department.

Table 16: Perceived employee benefits of LLM adoption for work.

Q7. Employee Benefit	Count	Share
Efficiency/Productivity	8	11.3%
Improve Writing/Grammar/Translation	6	8.5%
Information Search	6	8.5%
Structuring Documents	6	8.5%
Automate Repetitive Tasks	5	7.0%
Text Generation	5	7.0%
Email Composition	4	5.6%
Maintain/Improve Quality of Work	4	5.6%
Creativity Benefit	3	4.2%
Improve Programming	3	4.2%
More Focused/Interesting/Fun Work	3	4.2%
More Precise/Accurate	3	4.2%
Brainstorming/Inspiration	2	2.8%
Creating Presentations	2	2.8%

Good Starting Point	2	2.8%
Image/Video Generation	2	2.8%
AI/Personal Assistant	1	1.4%
Consistency	1	1.4%
Customer Service Improvement	1	1.4%
Getting a Concrete Answer	1	1.4%
(No Benefit)	1	1.4%
Provide Employee Independence	1	1.4%
Summarise Meetings/Calls	1	1.4%

5.5.2 Employee Perceived Risks

To understand the employee perspective on LLM adoption risks, the responses to Q8 were analysed and summarised in Table 17.

Q8. What do you perceive as the biggest risks to yourself when using LLMs for your own work?

The five most common perceived employee risks when it comes to using LLMs for work are inputting or exposing sensitive data via LLMs (14%), a lack of quality control or validation of LLM outputs (12%), developing an overreliance on LLMs (12%), model unreliability (10%), and bias present in LLM output or training data (6%).

Table 17: Perceived employee risks of LLM adoption for work.

Q8. Employee Risk	Count	Share
Inputting/Exposing Sensitive Data	7	13.7%
Lack of QC/Output Validation	6	11.8%
Overreliance on LLM	6	11.8%
Unreliability	5	9.8%
Bias	3	5.9%
Complacency	3	5.9%
Job Loss due to Automation	3	5.9%
Data Leakage/Breach	2	3.9%
Explainability/Black Box	2	3.9%
Hallucination/Delusion	2	3.9%
AI Aversion	1	2.0%
Ethical Risk	1	2.0%
IP Infringement	1	2.0%
Lack of Critical Thinking/Reflection	1	2.0%
Lack of Risk Awareness	1	2.0%
Loss of Tacit Knowledge/Expertise	1	2.0%
None	1	2.0%
Programming Risk	1	2.0%
Skill Deterioration	1	2.0%
Training Data Poisoning	1	2.0%
User Manipulation	1	2.0%
Work Becomes Dull/Uninspiring	1	2.0%

5.6 Organisational Perspective

This section analyzes the interview responses on organisations’ perceived benefits and risks associated with LLM adoption by way of a frequency analysis of the risk codes associated with Q9 and Q10.

5.6.1 Organisational Perceived Benefits

To understand the organisational perspective on LLM adoption benefits, the responses to Q9 were analysed and summarised in Table 18.

Q9. What do you perceive as the biggest benefits to your organisation when using LLMs in your work?

The five most common perceived organisational benefits when it comes to using LLMs for work are increased efficiency/productivity (28%), automating repetitive tasks (13%), gaining a competitive advantage (6%), cost savings (6%), and customer service improvement (6%).

Table 18: Perceived organisational benefits of LLM adoption for work.

Q9. Organizational Benefit	Count	Share
Efficiency/Productivity	13	27.7%
Automate Repetitive Tasks	6	12.8%
Competitive Advantage	3	6.4%
Cost Savings	3	6.4%
Customer Service Improvement	3	6.4%
Information Search	3	6.4%
Structuring Documents	2	4.3%
AI Democratization	1	2.1%
AI/Personal Assistant	1	2.1%
Amazement	1	2.1%
Data Quality Improvement	1	2.1%
Good Starting Point	1	2.1%
Improve Programming	1	2.1%
Improve Writing/Grammar/Translation	1	2.1%
Increased ROI	1	2.1%
More Human-to-Human Work	1	2.1%
More Precise/Accurate	1	2.1%
N/A	1	2.1%
New Opportunities	1	2.1%
Perform Analysis	1	2.1%
Scalability	1	2.1%

One of the mentions listed as ‘N/A’ comes from P12 who explained that although they “could imagine that LLMs could help the organisation,” they do not use LLMs in their work and therefore cannot think of any examples.

5.6.2 Organisational Perceived Risks

To understand the organisational perspective on LLM adoption risks, the responses to Q10 were analysed and summarised in Table 19.

Q10. What do you perceive as the biggest risks to your organisation when using LLMs in your work?

The five most common perceived organisational risks when it comes to using LLMs for work are the risk of a data leakage/breach (11%), a lack of quality control or validation of LLM outputs (11%), non-compliance with GenAI regulations (6%), bias present in LLM output or training data (6%), and a negative impact on customers (5%).

Table 19: Perceived organisational risks of LLM adoption for work.

Q10. Organizational Risk	Count	Share
Data Leakage/Breach	7	10.8%
Lack of QC/Output Validation	7	10.8%
Regulatory Compliance	4	6.2%
Bias	3	4.6%
Environmental Sustainability	3	4.6%
Negative Impact on Customers	3	4.6%
Reputational Risk	3	4.6%
Unreliability	3	4.6%
Widening Competitive Divide	3	4.6%
Data Security	2	3.1%
Inputting/Exposing Sensitive Data	2	3.1%
Lack of Transparency	2	3.1%
New Cyber Threats	2	3.1%
Overkill/Using Wrong Solution	2	3.1%
Poor Data Quality	2	3.1%
Bankruptcy	1	1.5%
Cultural Risk	1	1.5%
Explainability/Black Box	1	1.5%
Fraud	1	1.5%
Hallucination/Delusion	1	1.5%
Impersonal/Loss of Personal Touch	1	1.5%
Job Loss due to Automation	1	1.5%
Lack of Accountability	1	1.5%
Lack of Awareness	1	1.5%
Lack of Consistency	1	1.5%
Lack of Critical Thinking/Reflection	1	1.5%
Loss of Control	1	1.5%
Overreliance on LLM	1	1.5%
Phishing/Spam	1	1.5%
Ransom	1	1.5%
Single Point-of-Failure	1	1.5%
Unknown Risks	1	1.5%

5.7 Employee Expectations of LLMs

Understanding employees' expectations of LLMs is important for guiding organisations in their future adoption plans. These expectations are extracted from responses to Q13 on what employees are most excited about with regards to LLM adoption within their organisation. The responses are summarised in Table 20.

Q13. What are you most excited about with regards to LLM adoption within your organisation?

Table 20: Participant responses on what they are most excited about regarding LLM adoption (Q13).

Participant	Benefits
P01	Summarise Meetings/Calls
P02	Increased Understanding of LLMs
P03	Efficiency/Productivity; Creativity Benefit; Learning Opportunity; New Opportunities; More Focused/Interesting/Fun Work
P04	Efficiency/Productivity; Automate Repetitive Tasks
P05	Efficiency/Productivity; New Opportunities; Customer Service Improvement
P06	Summarise Documents
P07	Amazement; More Focused/Interesting/Fun Work
P08	Automate Repetitive Tasks;
P09	Efficiency/Productivity; Automate Repetitive Tasks; AI/Personal Assistant; More Focused/Interesting/Fun Work
P10	Improve Programming; Learning Opportunity; Getting a Concrete Answer; Creating Presentations
P11	Versatility
P12	Efficiency/Productivity
P13	N/A
P14	Efficiency/Productivity; More Focused/Interesting/Fun Work
P15	Optimism wrt LLMs; Amazement
P16	Optimism wrt LLMs; Cultural Benefit
P17	Efficiency/Productivity; Increased ROI
P18	New Opportunities; Customer Service Improvement; Data Quality Improvement

A frequency analysis of the perceived benefits (Table 21) reveals that participants are most excited about the potential efficiency and productivity gains that LLMs offer for their work (19%), how LLM usage could free up more time for focused, interesting, and fun work (11%), repetitive tasks becoming automated (8%), the new opportunities that LLMs present (8%) such as new business models, and improvements to customer service (5%). These benefits are inferred to be the employee expectations for the future use of LLM at their organisation and in their work.

Table 21: Perceived benefits of LLM that participants are most excited about, ordered by frequency.

Benefit	Theme	Frequency	Share
Efficiency/Productivity	Efficiency	7	18.9%
More Focused/Interesting/Fun Work	Creative	4	10.8%
Automate Repetitive Tasks	Efficiency	3	8.1%
New Opportunities	Business	3	8.1%
Amazement	Sentiment	2	5.4%
Customer Service Improvement	Business	2	5.4%
Learning Opportunity	Creative	2	5.4%
Optimism wrt LLMs	Sentiment	2	5.4%
AI/Personal Assistant	Assistance	1	2.7%
Creating Presentations	Creative	1	2.7%
Creativity Benefit	Creative	1	2.7%
Cultural Benefit	Sentiment	1	2.7%
Data Quality Improvement	Higher Quality	1	2.7%

Getting a Concrete Answer	Assistance	1	2.7%
Improve Programming	Efficiency	1	2.7%
Increased ROI	Business	1	2.7%
Increased Understanding of LLMs	Other	1	2.7%
Summarise Documents	Assistance	1	2.7%
Summarise Meetings/Calls	Assistance	1	2.7%
Versatility	Assistance	1	2.7%

These benefits have been associated to the same themes as in 5.1 Benefits Overview and a frequency analysis is shown in Table 22. The most common occurring themes are 'Creative' (20%) and 'Efficiency' (20%).

Table 22: Themes of perceived benefits that participants are most excited about, ordered by frequency.

Benefit Themes	Frequency	Share
Creative	4	20.0%
Efficiency	4	20.0%
Assistance	3	15.0%
Business	3	15.0%
Sentiment	3	15.0%
Higher Quality	1	5.0%
Other	1	5.0%
Literary	1	5.0%

6 Discussion

In this section, the previous research findings are brought together and adapted to TAM. First, an explanation of LLM's perceived ease of use (E) is provided while interview responses argue that employees perceive LLMs to be usefulness for their work (P). The consequence on the rest of the model of these two factors, P and E, being positive is elaborated upon with the objective of drawing conclusions on how organisations should best handle LLM adoption. A discussion on the validity and reliability of the research methods and results is addressed, as well as the study limitations inherent to LLM as a relatively new technology.

6.1 Technology Acceptance Model for LLM Adoption

LLMs have a high perceived ease of use. The most popular LLMs like ChatGPT allow users to interact with the model via a chatbot interface. This is an accessible way for employees to use the technology as it allows them to interact with it using natural language and does not necessitate programming knowledge to use. Moreover, the chat functionality means that the conversational context is stored in every chat session which allows users to iteratively interact with the model (OpenAI, 2022), for example, to ask clarifying questions. P10 specifically mentions the benefit of LLMs giving concrete answers, explaining that it is easier to use than Google for information search and is much more accessible than previous AI tools that his organisation has adopted. Indeed, easier information search was the third most mentioned perceived benefit of LLM adoption across the interview responses. Overall, the perceived ease of use of LLM is positively affected by the simplicity of natural language interaction through a chat interface that does not require users to have special knowledge or training to use. It should be mentioned, however, that knowledge and training on how to effectuate prompt engineer does improve the accuracy and quality of model outputs (Marvin et al., 2024).

Similarly, from the interviews, it is observed that employees highly appreciate LLMs for their capabilities to increase efficiency/productivity, improve writing/grammar/translation, search for information, structure documents, and automate repetitive tasks. This shows that there is a clear perceived usefulness of LLMs for

knowledge work, which in turn, positively affects employee attitudes towards using LLMs. However, there are a few important risks perceived by organizations, notably risks of sensitive data being leaked via LLMs and a lack of quality control or output validation by employees using model outputs in their work. These risks seem to be disproportionately concerning to organisations who risk reputational damage, while employees using the LLM are more concerned with the efficiency gains offered by these models for their work tasks. Herein lies a misalignment in the employee and organisational perspectives.

In lieu of these risks, organisations seek to control and mitigate them through usage policies to regulate how employees use LLMs in their work. Interview results found that all the usage policies described by participants make some mention about not inputting sensitive information into the model. Moreover, results on current adoption found that 2/9 organisations opted to ban LLM use at work which should eliminate LLMs' risks completely, while many others took a middle-of-the-road solution by allowing limited access to LLM, such as by request or only to specific departments/roles. These policies impose limits on the behavioral intention to use (B) LLMs by employees.

Despite these mitigating efforts by organisations, the intention to use is in some cases so strong that employees will actively violate policies against the use of LLMs. P16's interview revealed that Bank A's engineering team uses workarounds to access LLMs, showing that actual system use is likely higher than organisations are aware of. In fact, for some employees, the usage of such tools is so crucial to their work that they are choosing to leave their company for roles at competing firms that do authorize LLMs. This decision for programmers makes sense given previous research has found that 47-56% of worker tasks could be completed significantly faster at the same level of quality when incorporating software and tooling built on top of LLMs (Eloundou et al., 2023). The unavailability of LLM tools puts pressure on organisations to adopt LLMs as it poses a risk to the competitiveness of the firm (e.g., through the loss of talent) and creates a culture of division between the orders of top management and the reality of employee behaviour on the work floor.

Rather than prohibiting LLMs or restricting their usage, organisations are advised to find ways to incorporate them into their employee workflows by providing clear policies and guidelines. This approach will be most beneficial to creative and literary workflows, such as for improving writing, grammar, and translation, but also for information search, structuring documents, and text generation. Indeed, this trend is already reflected by major software vendors incorporating LLMs into all their core products (Cardon et al., 2023). Of note, Microsoft was found to be a leader in the current adoption of LLMs within the Dutch financial sector (accounting for over half of the LLMs used by participants) as the company incorporates their Copilot LLM into their core productivity and programming tools (Spataro, 2023).

To achieve this integration, organisations are advised to write and implement clear usage policies. This will ensure that the benefits of allowing LLM usage at work are enjoyed while mitigating the most important risks. Clear communication of items like which uses are and are not allowed, who is allowed to use it, and what kind of data may be inputted is crucial to avoid confusion. These specifications can be tailored to the needs of the organisation and according to their risk appetite.

Two options are proposed for choosing which LLM solution to adopt, depending on the organisation's resources and need for data security. Firstly, the LLM can be deployed locally to manage the data leakage risk. A privately run LLM ensures that all the data inputted by employees remains contained to the company environment. The drawback of this option is that such LLMs tend to be older models which are less powerful than the cloud-based state-of-the-art solutions from the largest software companies (e.g., Microsoft, Google, Meta, etc.). Additionally, if the LLM is hosted within the company's physical infrastructure, in-house expertise will be required to operate and maintain the model which can increase the overall cost. The second option is to hire a service providing company to implement it for the organisation. This is a lower cost of entry to LLMs which also allows access to the most powerful LLMs and does not incur any maintenance costs. However, it does mean that the data is sent to and stored with the third party. In this case, it is recommended that organisations find

an LLM provider capable of guaranteeing that sensitive data will not be collected or will be stored separately from other client data. In any case, this option has an inherently higher risk by virtue of the LLM being hosted in an external environment.

6.2 Validity & Reliability

To ensure the validity and reliability of the research findings, various measures were taken. Firstly, each interview followed a set structure including a questions list built around the research question and sub-questions. The interviews were conducted in a flexible and interactive way, giving participants the opportunity to expand on their answers and independently voice their experiences. Second, each transcript was meticulously reviewed to correct for errors in recording and to verify the accuracy of transcriptions and interpretations. Third, the research results inferred from the participant responses were cross-checked with existing LLM literature to confirm the validity of the claims and/or provide further explanation to a finding's motivations or context.

Given the novelty of LLMs, however, the people and roles responsible for advising on the adoption decision process are not consistent from one organisation to the other. A variety of roles was therefore studied, with each participant having different years of experience, educational background, and areas of expertise. It was noticed for instance that responses from risk officers tended to be more risk averse. Similarly, heads of AI were more knowledgeable about LLM technology given their expertise on the matter and tended to be more optimistic in their perceptions of LLMs than other roles. This variety in the participants reduces the validity of the results as compared to if all participants shared the same role.

In addition, P16's interview recording unfortunately only started halfway through the interview which meant that just a partial transcript was available for coding. Luckily, this error was spotted during the interview and a summary of the participants responses was discussed at the end of the interview to salvage the non-recorded answers. This means that although all of P16's answers were gathered, the first half of the questions were answer summaries which naturally affected how the transcript was coded as well as a loss of detail for the analysis.

On a few occasions, it was also noticed during the data analysis that certain answers were missing due to an inconsistency in questioning. This could be explained by human error from the interviewer such as accidentally skipping a question or sub-question, the interviewee's response going off on a tangent which meant that the question was not directly answered, and questions being skipped or not able to be answered due to running out of time. Although this happened exceptionally, it nevertheless reduces the reliability of the data collected.

Another limitation of this research is the coding of the interview transcripts. This was done manually by one person. Despite having undergone multiple revisions, the possibility that certain codes were missed always exists. There is also an inherent bias to having just one person perform coding, as they have a certain way of thinking about and formulating the codes. Naturally, some participants talked more than others, provided longer or more detailed answers, or repeated keywords more often which inflate the frequency of certain codes. A notable example of an inflated risk code was that of 'user manipulation', which was inflated in its presence by Q15 on the risk of user manipulation from LLMs which was ultimately not used.

Finally, the time gap between interviews is likely a significant factor influencing how participants responded to questions on current adoption and their organisation's usage policies, given how quickly LLM technology is being developed and adopted.

6.3 Study Limitations

LLM is a relatively new technology, and it cannot be assumed that all participants have a good understanding of what they are and how they work. Moreover, policies and regulations for LLMs are still being written or have

only recently been published. Organisations are contending with this novelty, also in their own formation and enforcement of usage policies. The fact that LLM is still the early stage of adoption means that things are changing quickly, so much so that the few weeks time difference between interviews is suspected to be a significant factor affecting the responses collected. An example hinting of this being the case was found in the discrepancy between participant responses to Pension Fund A's LLM usage policy. P04, one of the first participants to be interviewed, said that LLMs were not allowed, whereas three colleagues interviewed at later dates said that they were.

Similarly, AI roles are somewhat new in organisations, and there is no widespread set of AI roles expected for all organisations. Therefore, the distribution of roles interviewed is also uneven between organisations. For example, the heads of Data & AI interviewed worked exclusively in banking, whereas all the risk managers worked for insurance firms or pension funds. This could naturally influence the tendency of interview answers to be more optimistic vs risk aware on the perceptions related to LLM adoption. Similarly, four of the six pension fund interviewees were risk managers with the remaining two being privacy officers. This could (in part) motivate the more conservative adoption approach identified for pension funds.

Some participants had more expertise/experience with LLMs than others, while some showed hesitancy in their understanding of LLMs during the interview or implied this within their answers. P05's answer to the question on employee perceived benefit is one of 4 recorded instances of this doubt of knowledge: "For our department itself, I can't really think of things where artificial intelligence can help us with, but I think that's more my lack of knowledge." It is possible that participants' lack of LLM knowledge meant that they provided different perceptions based on a false understanding of how LLMs work.

7 Conclusions

7.1 Results Discussion

In the following sections, each research sub-question is answered by discussing the research findings.

7.1.1 Most Common LLM Use Cases

SQ1 can be answered from this research's findings on how LLMs are used by employees in practice as well as from the existing literature.

SQ1. What are the most common use cases of LLMs?

In this study, half of the participants (9/18) stated that they use LLMs in their work while the other half does not. Frequency analyses of which LLMs they use and what they use it for was performed. Firstly, it was found that the most common LLMs used are Microsoft's Copilot, OpenAI's ChatGPT, and private versions of ChatGPT. In total, more than half of the LLMs that participants use are Microsoft products which suggests that the company is responsible for a significant portion of the Dutch financial sector's LLM adoption thus far. This finding is also in agreement with the literature stating that major software vendors have already integrated LLM technology into all their core products (Cardon et al., 2023).

Secondly, the most common use cases stated by participants were found to be literary and creative in nature, and include creating presentations, generating text, composing emails, and structuring documents. The creation of presentations and the generation of text are creative tasks, composing emails requires social skills, and structuring documents requires analytical skills. Overall, the tasks require the use of theoretical knowledge to complete. This analysis of the use cases confirms that the participants' work is indeed knowledge work, "characterized by an emphasis on theoretical knowledge, creativity, and use of analytical and social skills (Frenkel et al., 1995, p. 773)."

The use cases found also show overlap with other literature on professional LLM usage. The study by Cardon et al. (2023) on GenAI in the workplace found that 42% of the US workers in their sample used ChatGPT to research a topic, 32% to draft an email or text, 26% to draft text for a longer document like a report, 22% to summarize text, and 21% to edit text. Once more, these use cases are creative and literary in nature and require theoretical knowledge to complete.

7.1.2 Employees' Perceived Benefits and Risks of LLM

In a bid to answer SQ2, frequency analyses were performed on interview responses about the perceived benefits and risks of LLM adoption for employees.

SQ2. What are employees' perceived benefits and risks associated with LLM adoption?

The results show that employees believe LLMs are most beneficial for their capabilities to increase efficiency/productivity, improve writing/grammar/translation, search for information, structure documents, and automate repetitive tasks. Increased efficiency/productivity is most frequently stated, accounting for 11.3% of all mentions. The frequency of this perceived employee benefit is further supported by the work of Dell'Acqua et al. (2023) which found that consultants using GPT4 for realistic consultant tasks were significantly more productive, completing 12.2% more tasks and 25.1% more quickly.

In contrast, employees' most important perceived risks of LLM usage include inputting or exposing sensitive data via LLMs, a lack of quality control or validation of LLM outputs, developing an overreliance on LLMs, bias present in LLM output or training data, and job loss due to LLM automation. All of these fit within the seven main risks clusters of GenAI that were submitted in the literature review by Wach et al. (2023).

7.1.3 Organisations' Perceived Benefits and Risks of LLM

In a bid to answer SQ3, frequency analyses were performed on interview responses to the perceived benefits and risks of LLM adoption for organisations.

SQ3. What are organisations' perceived benefits and risks associated with LLM adoption?

The results showed that organisations stand to benefit most from LLM usage thanks to increased efficiency/productivity, automating repetitive tasks, gaining a competitive advantage, cost savings, and customer service improvement. In contrast, the most important perceived risks of LLM usage for organisations include the risk of a data leakage/breach, a lack of quality control or validation of LLM outputs, non-compliance with GenAI regulations, bias present in LLM output or training data, and a negative impact on customers.

The organizational and employee perspectives can be compared. Of the top-five perceived benefits for each one, increased efficiency/productivity and automating repetitive tasks are shared benefits. Indeed, the perceived benefit of LLMs providing increased efficiency/productivity ranks highest for both perspectives in addition to being the most frequently mentioned benefit across all interview questions.

The other three top-five organisational perceived benefits (gaining a competitive advantage, cost savings, and customer service improvement) were not at all mentioned as a perceived benefit for employees. This makes sense seeing as these benefits are directly related to helping the business, providing little to no benefit to the mentioned use cases of the interviewed employees.

A lack of quality control/output validation and bias are common perceived risks between the two perspectives. There is an additional commonality to be found between the top-five perceived risks which are closely related yet distinctly different, namely, the top employee perceived risk of inputting/exposing sensitive data and the top organisational perceived risk of data leakage/breach. Both are cybersecurity risks related to failing to safeguard sensitive or private data. However, the inputting/exposure of sensitive data is likely more concerning to an employee who is the user inputting information into the LLM, whereas a data leakage/breach is something that occurs to the organisation rather than affecting single individuals.

7.1.4 Future Adoption and Employee Expectations

Lastly, to answer SQ4, responses on organisations' future adoption plans were analysed and compared to employee expectations of LLMs derived from what employees are most excited about with regards to LLM adoption.

SQ4. What are organisations' current levels of LLM adoption and their future LLM adoption plans, and how are they aligned with employee expectations of LLMs?

In terms of organisations' current levels of LLM adoption, the interviews found that that only a minority of financial organisations prohibit employees from using LLMs for their work due to not having a good enough understanding of LLMs, not having the proper safeguards in place for its use at work, the risk of exposing sensitive data that is inputted in the LLM, and the organisation's lower risk appetite. 14/18 participants said that their organisation allows LLMs to be used at work, 3 of which allow it with limited use, and 4/18 said that they are prohibited. In terms of future LLM adoption plans, 13/18 participants were aware of their organisation's plans, 3 were unclear, and 2 said that their organisation does not (yet) have plans. Most organisations (14/18) have already taken the first step of adoption by running one or several LLM pilots to test the technology and its usefulness.

Current usage policies can be categorised in three ways: a complete ban on LLMs to eliminate the risks of LLM usage, or at least temporarily until new LLM policies are finalised and implemented; a partial ban which is the most common type, whereby certain LLMs are prohibited but not others, or whereby LLMs are allowed to

select roles/departments or use cases (e.g., pilot development); free use which is the least common type, whereby firms trust their employees to use LLMs responsibly and rely on their common sense. Most organisations have also already implemented pilots to assess future use cases of LLMs prior to adoption and research their added value. An emphasis is placed on risk assessments to make sure the risks are acceptable before continuing and so that the necessary safeguards can be implemented.

In terms of future LLM adoption plans, the targeted capabilities described by organisations that already know what they would like to achieve with their future adoption plans are summarised as follows:

- Using Copilot to help programmers write better code.
- Using GPT to analyse conversations between customers and the help desk staff by transcribing data speech-to-text and making summaries, gathering new insights from this data.
- Using GPT to analyse emails and predict customers' questions.
- Using Bing as a general chatbot assistant to answer employees' questions.
- Using GPT on internal knowledge models by asking the LLM a question for it to return the correct documents from the company database.

When it comes to employee expectations, participants are most excited about the potential efficiency and productivity gains that LLMs offer for their work (19%), how LLM usage could free up more time for focused, interesting, and fun work (11%), repetitive tasks becoming automated by LLM (8%), the new opportunities that LLMs present (8%) such as new business models, and improvements to customer service (5%).

There is general overlap between the targeted capabilities and employee expectations, such as the productivity gains from helping programmers write better code, or improvements to customer service as from using GPT to analyse emails and predict customer questions. Gathering new insights from customer help desk conversations can also be viewed as a new opportunity provided by LLMs, and having a chatbot to answer employee questions frees up time and focus away from menial tasks and towards more focused, interesting, and fun work. This overlap shows that there is a general alignment between what employees expect to be able to do with LLMs in the future and their organization's future LLM adoption plans, answering the second half of SQ4.

7.2 Research Conclusions

Using a qualitative research approach, this study was aimed at exploring employee and organisational perceptions on the benefits and risks of LLM adoption within the Dutch financial sector. This study has been conducted at a global Professional Services Firm and used their existing network of people and clients within the Dutch financial sector to conduct a total of 18 open-ended interviews with experts who act as advisors to top management in the decision-making process of new technology adoption such as LLMs. The research question to be answered is:

- How do employees' and organisations' perceived benefits and risks of Large Language Models (LLMs) influence financial organisations' LLM adoption plans?

To answer this question, TAM was used as a theoretical framework to bring together the interview findings. Applying the model to LLM adoption showed that there is both a high perceived ease of use as LLMs are often interfaced through chatbots in natural language, and a high perceived usefulness as they improve employees' ability to achieve their most common work tasks quicker and more efficiently. This perceived usefulness of LLMs is posited to positively affect employee attitudes towards using LLMs. However, there are a few important risks perceived by organizations, notably risks of sensitive data being leaked via LLMs and a lack of quality control or output validation by employees using model outputs in their work. These risks seem to be disproportionately concerning to organisations who risk reputational damage, while employees using the LLM are more concerned with the efficiency gains offered by these models for their work tasks, presenting a misalignment in the employee and organisational perspectives.

Instead of restricting LLM usage, it is recommended that organisations find ways to incorporate them into their employee workflows by providing clear policies and guidelines. It was found that this approach will be most beneficial to creative and literary workflows like improving writing/grammar/translation, information search, structuring documents, and text generation.

To achieve this integration, organisations should write and implement clear usage policies. This will ensure that the benefits of allowing LLM usage at work are enjoyed while mitigating the most important risks. Clear communication of items like which uses are and are not allowed, who is allowed to use it, and what kind of data may be inputted is crucial to avoid confusion. These specifications can be tailored to the needs of the organisation and according to their risk appetite.

7.3 Future Work

The open-ended approach taken during this study allowed for a broader capture of perspectives and paves the way for future research to explore in more depth how particular factors influence the LLM adoption decision-making process. Future work can be done to investigate particular themes of benefits and risks or reproduce the study within another business sector. Furthermore, it would be interesting to compare companies who adopted LLMs to those who did not, and quantitatively evaluate their development.

Additionally, future work could choose to analyse how the perceptions surrounding LLM influence its adoption by means of a different framework than TAM like adoption theory, or by using emerging theory on different dimensions of analysis to the employee vs organisational perspectives, such as organisation size (e.g., large vs small organisations), industry (e.g., comparing organisations across sectors), or type of work (e.g., knowledge vs non-knowledge work).

Appendix A: Employee and Organisational Perceived Benefits and Risks of LLM Adoption

Table 23: Interview responses (coded) on employee and organisational perceived benefits/risks (Q7-Q10).

Participant	Q7. User Benefits	Q8. User Risks	Q9. Org Benefits	Q10. Org Risks
P01	Creativity Benefit; Structuring Documents; Improve Writing/Grammar/Translation; Email Composition	Lack of QC/Output Validation	Efficiency; Structuring Documents; Good Starting Point	Lack of Consistency; Lack of QC/Output Validation; Overreliance on LLM;
P02	Automate Repetitive Tasks; Summarise Documents; Creativity Benefit; Improve Writing/Grammar/Translation; Improve Programming; Email	AI Aversion; User Manipulation	Automate Repetitive Tasks	Explainability/Black Box; Hallucination/Delusion; Data Leakage
P03	Efficiency/Productivity; Summarise Meetings/Calls; Information Search; Maintain/Improve Quality of Work; More Focused/Interesting/Fun Work;	Unreliability; Training Data Poisoning; Lack of QC/Output Validation; Inputting/Exposing Sensitive Data; Overreliance on LLM;	Efficiency/Productivity	Unknown Risks; Reputational Risk; Lack of QC/Output Validation
P04	Improve Writing/Grammar/Translation; Text Generation; Information Search; Efficiency/Productivity.	Overreliance on LLM; Work Becomes Dull/Uninspiring; Inputting/Exposing Sensitive Data; Job Loss due to Automation	Efficiency/Productivity; More Precise/Accurate; Cost Savings; Structuring Documents; Competitive Advantage.	Lack of Accountability; Lack of Transparency; Data Leakage/Breach; New Cyber Threats; Bias
P05	Information Search; Customer Service Improvement; Automate Repetitive Tasks.	Programming Risk; Hallucination/Delusion; Unreliability	Efficiency/Productivity; Information Search	Reputational Risk; Negative Impact on Customers; Data Leakage/Breach; Fraud; Regulatory Compliance; Widening Competitive Divide
P06	None	Data Leakage/Breach; Inputting/Exposing Sensitive Data	Information Search; Efficiency/Productivity;	Lack of QC/Output Validation; Unreliability
P07	Summarise Documents; Automate Repetitive Tasks; Maintain/Improve Quality of Work; Good Starting Point; Text Generation; Efficiency/Productivity; More Focused/Interesting/Fun Work.	Overreliance on LLM; IP Infringement; Inputting/Exposing Sensitive Data; Job Loss due to Automation	Cost Savings; Efficiency/Productivity; Customer Service Improvement	Negative Impact on Customers; Data Leakage/Breach; Inputting/Exposing Sensitive Data; Regulatory Compliance; Sustainability Concerns.
P08	Automate Repetitive Tasks; More Precise/Accurate	Unreliability; Hallucination/Delusion	Automate Repetitive Tasks	Regulatory Compliance; Data Leakage/Breach
P09	Information Search; Efficiency/Productivity; Summarise Documents; Improve Writing/Grammar/Translation; Email Composition; AI/Personal Assistant; Text Generation; Improve Programming; Information Search	Inputting/Exposing Sensitive Data; Data Leakage/Breach; Explainability/Black Box; Bias; Overreliance on LLM; Lack of QC/Output Validation	Competitive Advantage; Efficiency/Productivity; Data Quality Improvement	Bankruptcy; Regulatory Compliance; Widening Competitive Divide
P10	Efficiency/Productivity; Improve Programming; Text Generation;	None	Efficiency/Productivity; AI/Personal Assistant; Improve Writing/Grammar/Translation	Poor Data Quality; Sustainability Concerns; Overkill/Using Wrong Solution; Loss of Control.
P11	More Focused/Interesting/Fun Work	Ethical Risk; Job Loss due to Automation	More Human-to-Human Work; Automate Repetitive Tasks; Cost Savings	
P12	Getting a Concrete Answer	Explainability/Black Box; Inputting/Exposing Sensitive Data; Lack of Risk Awareness	N/A	Unreliability; Poor Data Quality
P13	Efficiency/Productivity; Consistency; Maintain/Improve Quality of Work	Complacency;	Efficiency/Productivity; Competitive Advantage	Lack of Awareness; Negative Impact on Customers; Overkill/Using Wrong Solution; Lack of QC/Output Validation.
P14	Brainstorming/Inspiration; Information Search; More Precise/Accurate; Improve Writing/Grammar/Translation; Good Starting Point	Lack of QC/Output Validation; Misinformation/False Positives/Negatives	Automate Repetitive Tasks; Perform Analysis	Inputting/Exposing Sensitive Data; Data Leakage/Breach; Bias
P15	More Precise/Accurate; Summarise Documents; Efficiency/Productivity; Creating Presentations; Brainstorming/Inspiration	Lack of QC/Output Validation; Loss of Tacit Knowledge/Expertise; Overreliance on LLM	Efficiency/Productivity; Automate Repetitive Tasks	Sustainability Concerns; Job Loss due to Automation
P16	Maintain/Improve Quality of Work	Complacency; Lack of Critical Thinking/Reflection	Customer Service Improvement; Improve Programming;	Lack of QC/Output Validation; Lack of Critical Thinking/Reflection.
P17	Automating Repetitive Tasks; Improve Writing/Grammar/Translation	Inputting/Exposing Sensitive Data; Bias; Misinformation/False Positives/Negatives; Skill Deterioration	Scalability; Efficiency/Productivity; Automate Repetitive Tasks; Increased ROI; Customer Service Improvement	Lack of QC/Output Validation; Misinformation/False Positives/Negatives; Bias; Data Security; Single Point of Failure
P18	Efficiency/Productivity; Email Composition; Text Generation; Structuring Documents; Creating Presentations; Provide Employee Independence; Image/Video Generation; Creativity Benefit.	Lack of QC/Output Validation; Overreliance on LLM; Complacency; Bias.	Efficiency/Productivity; Information Search; Amazement; AI Democratization; New Opportunities	Impersonal/Loss of Personal Touch; Lack of Transparency; Reputational Risk; Lack of QC/Output Validation; Cultural Risk; Widening Competitive Divide; Data Leakage/Breach; New Cyber Threats; Data Security; Phishing/Spam; Ransom

Appendix B: Summary of Future LLM Adoption Plans

Table 24: Summary of LLM future adoption plans including organizations' use of pilots.

Organization	Future Plans?	Pilot?	Participant	Q6. Future Adoption Plans? (Summarized)
Bank A (large)	Yes	Yes	P02	Yes. The bank is running a couple of LLM pilots. One is on Copilot helping coders write better code. Not for any code that gets put into production, so everything is to be checked and validated by humans of course. But the bank is experimenting with a couple of use cases internally.
	Yes	Yes	P16	Yes: "There's no official approval yet that the pilot goes to production, but we are still working on it. We do have some projects ongoing, but these are very sensitive and we only go into production after we have the policies and measures in place to do that." // once it is available, will be made available to all employees, not just engineers
Bank B (large)	Yes	Yes	P03	Yes. Bank B is doing the piloting to determine what we want to do with them. So, the intention is yes, if they're useful enough and the risks can be managed properly, then certainly there is a drive to use them more.
Bank C (large)	Unclear	Yes	P10	Unclear: "Now we're using Copilot. I wouldn't be surprised if we would have more in the future."
	Unclear	Yes	P14	Unclear: "Yes. I think there are different pilots, but I'm not very aware of them. I know there are definitely teams working on this, but the exact details I have no idea."
Bank D (small)	Yes	Yes	P17	"Yes, definitely. So, we have one [colleague] who is focusing on LLMs and making sure that we have the state-of-the-art available. That's it. So, whenever there's something better, we test it and use it."
Bank E (small)	Yes	Yes	P18	Yes: "In the future, we will have managed access to LLMs (including ChatGPT). Then we will give you a pop-up that says, hey, we see that you're trying to upload this data, or naming this, or inputting this. Reconsider. We're heavily investing with our teams to increase the number of solutions in our company that have LLMs incorporated. It doesn't mean they're solely based on LLMs but that they are part of the solution."
Insurance A	Yes	No	P01	Yes. We have plans to use other, more specific LLMs, and what we are doing right now is that we create solutions based on the technology of ChatGPT for very specific purposes within our own company.
	Yes	No	P13	Yes: "Within [department] we don't have so many models yet. We are designing the first use case to make the work of security documentation easier. So, we want to explore the possibilities in our own field."
Insurance B	Yes	Yes	P07	Yes. Insurance B is developing the use of GPT for analysing customer help desk conversations. The firm is also looking to use LLMs to analyse emails and predict customer questions. The firm will adopt LLM applications from Microsoft.
	Yes	Yes	P11	Yes. Insurance B has one application where they try to predict the chance that clients will contact them for different reasons. This one is to optimise our service. And there's some experiments with speech-to-text to analyse client conversations and try to gather more insights in the data that we can base our decisions on.
	Yes	Yes	P15	Yes: 1) "we transcribe data speech-to-text and then we use models to make summaries. And we also get new insights from that data. So, we aggregate it, make sure that we correlate the high over Google Classifications or NPSC classifications, or different things, depending on the business line." 2) "Another use case is of course internal knowledge models. So, you have your own database with knowledge or it's a semi structured data or it's PDFs or things that you need for policy documents. You can ask a question about it and it pops up the correct documents that contain that information." // "There are quite a lot of use cases at the moment, so I'm now talking about just two of the use cases, but I think we have 10 or 12 at different stages of development in different parts of the organisation."
Pension Fund A	Yes	Yes	P04	Yes. Pension Fund A is working on allowing LLMs in a controlled environment. Not a private LLM but rather via APIs to access an external LLM.
	Yes	Yes	P08	Yes. There are a couple of pilots including some talk about a Bing pilot (chatbot)
	Yes	Yes	P09	Yes.
	Unclear	Yes	P12	Unclear: "I would say yes. But you never know what the future brings us."
Pension Fund B	No	No	P05	No. Pension Fund B is not that far yet.
	No	No	P06	No. Pension Fund B is not that far yet.

Bibliography

- Campbell, S., Greenwood, M., Prior, S., Shearer, T., Walkem, K., Young, S., Bywaters, D., & Walker, K. (2020). Purposive sampling: Complex or simple? Research case examples. *Journal of Research in Nursing*, 25(8), 652–661. <https://doi.org/10.1177/1744987120927206>
- Cardon, P. W., Getchell, K., Carradini, S., Fleischmann, C., & Stapp, J. (2023). *Generative AI in the Workplace: Employee Perspectives of ChatGPT Benefits and Organizational Policies* [Preprint]. SocArXiv. <https://doi.org/10.31235/osf.io/b3ezy>
- Coyne, R. (2005). Wicked problems revisited. *Design Studies*, 26(1), 5–17. <https://doi.org/10.1016/j.destud.2004.06.005>
- Davis, F. (1987). *User Acceptance of Information Systems: The Technology Acceptance Model (TAM)*. University of Michigan.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality* (SSRN Scholarly Paper 4573321). <https://doi.org/10.2139/ssrn.4573321>
- Demystifying the Technology Acceptance Model (TAM): A Comprehensive Guide*. (2024, May 1). UserSense. <https://www.usersense.io/knowledge-base/usability-metrics/technology-acceptance-model-tam>
- Diasio, D. (2023, May 24). *Why enabling AI’s full value requires top-down thinking*. Ey.Com. https://www.ey.com/en_gl/consulting/why-enabling-ais-full-value-requires-top-down-thinking
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models* (arXiv:2303.10130). arXiv. <https://doi.org/10.48550/arXiv.2303.10130>
- EU AI Act: First regulation on artificial intelligence*. (2023, June 8). Europa.Eu. <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- Floridi, L. (2023). *AI as Agency Without Intelligence: On ChatGPT, Large Language Models, and Other Generative Models* (SSRN Scholarly Paper 4358789). <https://doi.org/10.2139/ssrn.4358789>

- Frenkel, S., Korczynski, M., Donoghue, L., & Shire, K. (1995). Re-Constituting Work: Trends towards Knowledge Work and Info-Normative Control. *Work, Employment and Society*, 9(4), 773–796.
<https://doi.org/10.1177/095001709594008>
- Goldman Sachs. (2023, November 29). *AI investment forecast to approach \$200 billion globally by 2025*. Goldman Sachs. <https://www.goldmansachs.com/intelligence/pages/ai-investment-forecast-to-approach-200-billion-globally-by-2025.html>
- Graham, B. (2023, February 1). *70-percent-of-workers-using-chatgpt-at-work-are-not-telling-their-boss*. Fishbowl Insights. <https://www.fishbowlapp.com/insights/70-percent-of-workers-using-chatgpt-at-work-are-not-telling-their-boss/>
- Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base—Analyst note. *Reuters*.
<https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Jetha, A., Shamaee, A., Bonaccio, S., Gignac, M. A. M., Tucker, L. B., Tompa, E., Bültmann, U., Norman, C. D., Banks, C. G., & Smith, P. M. (2021). Fragmentation in the future of work: A horizon scan examining the impact of the changing nature of work on workers experiencing vulnerability. *American Journal of Industrial Medicine*, 64(8), 649–666. <https://doi.org/10.1002/ajim.23262>
- Klarna. (2024, February 27). *Klarna AI assistant handles two-thirds of customer service chats in its first month*. <https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/>
- Klenk, M. (2023). *Ethics of Generative AI and Manipulation: A Design-Oriented Research Agenda* (SSRN Scholarly Paper 4478397). <https://doi.org/10.2139/ssrn.4478397>
- Klenk, M., & Hancock, J. T. (2019, December 19). *Autonomy and online manipulation*. Internet Policy Review. <https://policyreview.info/articles/news/autonomy-and-online-manipulation/1431>
- Madiega, T. (2023). *Artificial intelligence act*. European Parliament.
- Marvin, G., Hellen, N., Jjingo, D., & Nakatumba-Nabende, J. (2024). Prompt Engineering in Large Language Models. In I. J. Jacob, S. Piramuthu, & P. Falkowski-Gilski (Eds.), *Data Intelligence and Cognitive Informatics* (pp. 387–402). Springer Nature. https://doi.org/10.1007/978-981-99-7962-2_30
- MSc Management of Technology*. (2024, April 19). Tudelft.NL.
<https://www.tudelft.nl/en/education/programmes/masters/mot/mot>

- Navarra, K. (2023, March 3). *Using ChatGPT Correctly on the Job*. Society for Human Resource Management. <https://www.shrm.org/topics-tools/news/technology/using-chatgpt-correctly-job>
- Newell, S., Scarbrough, H., & Swan, J. (2009). *Managing Knowledge Work and Innovation*. Bloomsbury Publishing.
- OpenAI. (2022, November 30). *Introducing ChatGPT* [Blog]. <https://openai.com/blog/chatgpt>
- OWASP. (2024, January 6). *About the OWASP Foundation | OWASP Foundation*. <https://owasp.org/about/>
- Pyöriä, P. (2005). The concept of knowledge work revisited. *Journal of Knowledge Management*, 9(3), 116–127. <https://doi.org/10.1108/13673270510602818>
- Ray, S. (2023, May 19). Apple Joins A Growing List Of Companies Cracking Down On Use Of ChatGPT By Staffers—Here’s Why. *Forbes*. <https://www.forbes.com/sites/siladityaray/2023/05/19/apple-joins-a-growing-list-of-companies-cracking-down-on-use-of-chatgpt-by-staffers-heres-why/>
- Spataro, J. (2023, September 21). *Announcing Microsoft 365 Copilot general availability and Microsoft 365 Chat*. Microsoft 365 Blog. <https://www.microsoft.com/en-us/microsoft-365/blog/2023/09/21/announcing-microsoft-365-copilot-general-availability-and-microsoft-365-chat/>
- Wach, K., Duong, C. D., Ejdys, J., Kazlauskaitė, R., Korzynski, P., Mazurek, G., Paliszkiwicz, J., & Ziemba, E. (2023). The dark side of generative artificial intelligence: A critical analysis of controversies and risks of ChatGPT. *Entrepreneurial Business and Economics Review*, 11(2), 7–30. <https://doi.org/10.15678/EBER.2023.110201>
- Walkowiak, E., & MacDonald, T. (2023). *Generative AI and the Workforce: What Are the Risks?* (SSRN Scholarly Paper 4568684). <https://doi.org/10.2139/ssrn.4568684>
- Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., ... Gabriel, I. (2022). Taxonomy of Risks posed by Language Models. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 214–229. <https://doi.org/10.1145/3531146.3533088>
- Weiss, D. . C. (2023, March 16). *Latest version of ChatGPT aces bar exam with score nearing 90th percentile*. ABA Journal. <https://www.abajournal.com/web/article/latest-version-of-chatgpt-aces-the-bar-exam-with-score-in-90th-percentile>
- Wilson, S., & Dawson, A. (2023). *OWASP Top 10 for LLM Applications v1.1*. OWASP. <https://LLMtop10.com>

Zohny, H., McMillan, J., & King, M. (2023). Ethics of generative AI. *Journal of Medical Ethics*, 49(2), 79–80.

<https://doi.org/10.1136/jme-2023-108909>