



Universiteit
Leiden
The Netherlands

Opleiding Informatica & Economie

Bottom-Up Adoption, Reactive Governance: A Qualitative Study of Individual and Organisational Forces Shaping SME Generative AI Adoption

Jakub Korus

First supervisor:
T.D. Offerman MSc

Second supervisor:
Dr. N. Amat-Lefort

BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)
www.liacs.leidenuniv.nl

10 /07/2026

Abstract

Generative AI has been adopted at an exceptional pace, with over 88% of enterprises having adopted it by 2026. This is merely three years after the launch of its first commercially viable product, ChatGPT. This study explores how small to medium enterprises (SMEs) adopt generative AI: which tools are adopted for which use cases, and which enabling or inhibiting factors influence adoption. Previous literature has identified numerous organisational and various individual factors conditioning adoption, with some of the former shaping the latter. However, the two perspectives are rarely unified, despite evidence that they are interrelated. The SME context, moreover, is comparatively underserved. Even though its resource-constrained environment may serve as an incubator for innovation and creative governance.

This qualitative study employed an exploratory design in which semi-structured interviews were conducted with 19 SME employees from the Netherlands, Belgium and the United Kingdom. The companies in which these individuals were employed spanned numerous industries. Data were analysed through inductive coding, from which five themes emerged. It finds that generative AI adoption in SMEs is bottom-up and individual-led, spreading through informal conversations, peer observation and internal experts, with firms reacting to rather than proactively shaping adoption. Formal policy is an exception and unwritten norms the rule. These rules emerge as a response to risk, such as data privacy, reputational but also output quality and loss of control over autonomous generative AI. Adopters generally favour general-purpose tools over specialised ones, though specific motivations lead the latter to be adopted in particular cases.

This research primarily suggests that in the SME setting the individual and organisation should not only be studied apart, as each shapes the other. Adoption is the product of their interaction. Secondly, given the rapid and seemingly endless improvements in generative AI's capabilities, contemporary governance strategies in SMEs may be insufficient to effectively shape individual adoption of generative AI in the future. This raises the broader question of how organisations should design governance for technologies that are adopted bottom-up and develop rapidly, balancing enabling and restricting their use.

Acknowledgements

First of all, I would like to deeply thank Tyron Offerman for his continuous support during this project. His guidance and input as a supervisor during our frequent meetings were priceless in improving this study's quality, and, beyond that, in developing me as an individual. I am grateful that we embarked on this journey together.

My gratitude likewise goes to Natalia Amat-Lefort for her support and useful feedback on this thesis. Your probing during our initial meeting helped me to understand my own research gap better, which improved my own grasp on what I intended to do.

I would like to thank Reinout Offringa for his assistance in obtaining interviews.

Above all else, I am grateful for all individuals who chose to participate in the interviews. This project would not be possible without you.

And, of course, my mom and dad, for their endless love and for patiently listening to me ramble on about generative AI adoption in the past few months.

Contents

1	Introduction	1
2	Related Works	4
2.1	Concepts of Artificial Intelligence	4
2.2	Emergence of Generative AI	5
2.3	State of Practice in Generative AI	6
2.4	Theoretical Foundations of Technology Adoption	8
2.4.1	Individual frameworks	9
2.4.2	Organisational frameworks	10
2.5	Generative AI adoption	11
2.5.1	Individual Perspectives	11
2.5.2	Organisational Perspectives	11
2.6	Research Gap	12
3	Methodology	14
3.1	Research Context	14
3.2	Research Design	14
3.3	Sampling	15
3.4	Data Collection	15
3.5	Data Analysis	15
3.6	Coding Scheme	16
4	Results	18
4.1	Sample Overview	18
4.2	Tools in Use	19
4.3	Use Cases	21
4.4	Factors Shaping Adoption	22
4.4.1	Origins and Diffusion of Adoption	23
4.4.2	The Human Factor of Adoption	25
4.4.3	Tool Provision and Its Workarounds	26
4.4.4	Rules that Help and Hinder	28
4.4.5	Protecting Data and Judgement	31
4.5	Variation across Organisational Size	34
5	Discussion	36
5.1	Summary of Findings	36
5.2	Interpretation	37

5.3	Theoretical Contributions	39
5.4	Implications for Practice	40
5.5	Limitations	40
5.6	Future Research	41
6	Conclusions	42
	References	43
A	Consent Form	47
B	Interview Protocol	48
C	Proposed Survey Instrument	49

Chapter 1

Introduction

Since its inception in 2017, generative AI has been adopted at an unprecedented pace. Its adoption has outpaced both that of the internet and the personal computer, by a substantial margin (Bick, Blandin, and Deming 2026). ChatGPT, one of the top generative AI tools, broke records. It took TikTok nine months and Instagram 2.5 years to reach 100 million monthly active users, whereas ChatGPT did it in approximately three months (Hu 2023). The technology has not escaped industries' attention either, with over 88% of enterprises having adopted it to date (Stanford Institute for Human-Centered Artificial Intelligence 2026). Fig. 1.1 displays the adoption rates of the internet, the computer and Generative AI since their first mass-market product. There are established trend lines for the computer and internet as these have an extensive history. Generative AI, as shown in the figure under AI, has relatively only been recently introduced. Therefore, only one the single blue point represents its adoption rate.




Figure 1.1. Generative AI adoption rate alongside to the internet and computer adoption rates, adapted from (Bick, Blandin, and Deming 2026)

The technology has not escaped academic attention. This has prompted numerous studies of generative AI adoption across varied contexts and through multiple theoretical frameworks. Adoption research divides into two major streams: individual and organisational adoption. Across the two streams, three prominent models have been employed to study generative AI adoption. First, the Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh, Morris, et al. 2003) was used to explain the variation in technology use among individuals. This framework has four predictors—Performance Expectancy (PE), Effort Expectancy (EE), Social Influence (SI), and Facilitating Conditions (FC)—all of which influence Behavioural Intent (BI) to adopt a technology. To study organisational adoption, both the Technology–Organisation–Environment (TOE) framework (Tornatzky and Fleischer 1990), which divides factors into its three categories, and Institutional Theory (DiMaggio and Powell 1983), which describes the organisation being shaped by coercive, mimetic, and normative pressures, were frequently selected.

At the individual level, PE, SI and EE arose as the core influences on BI. However, their influence is context-dependent, varying with adoption stage, culture, familiarity and organisational rank (Kim, Blazquez, and Oh 2024; Fang et al. 2026; Xin, C. Zhang, and Peng 2025; Biloš and Budimir 2024). In particular, SI appears stronger in collectivistic settings, such as China, than in Western European ones (Fang et al. 2026). At the organisational level, relative advantage, top management support, internal skills, and competitive pressure are recurring drivers. Institutional pressures are ambivalent, where they enable and, under uncertainty, dampen adoption (L. Zhang et al. 2025). The regulatory environment remains contested. Complexity and immaturity of tools acted as a direct barrier.

Problem Statement

Previous literature has investigated generative AI adoption within the SME context by examining individual and organisational adoption separately. However, evidence points to the two being intertwined; organisational hierarchy moderates SI (Xin, C. Zhang, and Peng 2025), supervisory support drives SI (Kim, Blazquez, and Oh 2024) and task-tool fit shapes PE (Xia and Chen 2025). Despite that, both levels are rarely united to bring a unified perspective of generative AI adoption (Y. M. Yesuf and Ziska Fields 2025). Lim et al. (2025) and Kumar et al. (2025) explicitly call for greater attention to the organisational contexts conditioning individual adoption. The SME context, in particular, is relevant as it is comparatively less studied while exhibiting distinct dynamics. This is a context in which organisational decisions are frequently shaped by individual ones. Together, the separation of the two levels and the under-studied SME context leave an underdeveloped perspective on how generative AI is adopted in small to medium enterprises.

Research Questions

This thesis aims to uncover how organisational and individual factors influence each other, and together affect generative AI adoption in small to medium enterprises. It does so by answering the following question:

RQ: *How do SMEs overcome barriers and create enablers to increase generative AI adoption?*

- **SQ1:** Which specific generative AI tools (e.g., foundation model API calls, ChatGPT, Microsoft Copilot) are adopted in SMEs?
- **SQ2:** For which functional business use-cases are these tools applied in SMEs?
- **SQ3:** How are organisational barriers and enablers shaping individual adoption of generative AI in SMEs?

Thesis Outline

The remainder of this thesis is structured as follows. The next chapter reviews related literature covering how generative AI fits into the broader Artificial Intelligence landscape and how it developed. It also covers the technology's application in industry, the theoretical frameworks of information systems adoption research and how these frameworks have been applied to investigate generative AI adoption.

The subsequent methodology chapter describes how this study was designed. In broad strokes, it covers who could be interviewed, how data would be collected and later analysed. Importantly, the chapter introduces crucial definitions necessary to understand the findings.

The results chapter follows by first presenting the obtained sample. Further, it examines the tools adopted, the use cases for which these tools have been applied, and factors that influenced their adoption. It concludes by showing how tool use and factors vary as partitioned by organisational size.

The discussion interprets these findings in relation to existing literature. Then, it sets out theoretical and practical contributions, and confesses the study's limitations and future research directions. Finally, the conclusion summarises and closes the study.

Chapter 2

Related Works

2.1 Concepts of Artificial Intelligence

The field of Artificial Intelligence (AI) is commonly defined as the study of machines capable of performing tasks that require intelligence (Russell and Norvig 2010). It has numerous subfields, of which Machine Learning (ML) is of particular interest. It is the study of algorithms that enable computers to learn patterns from data and improve their performance on tasks without being explicitly programmed (Mitchell 1997). Machine learning methods are commonly applied to tasks such as regression, which involves predicting numerical values, and classification (Krizhevsky, Sutskever, and G. Hinton 2012), which involves assigning categorical labels to observations, among others. They have demonstrated strong empirical performance across the previous prediction tasks when trained on large enough datasets (Halevy, Norvig, and Pereira 2009). However, limitations arise when prediction tasks require more advanced representations to be learned that the methods cannot adequately perform (LeCun, Bengio, and G. Hinton 2015).

A potential remedy to these limitations is provided by neural networks (NNs). Neural networks are computational models composed of interconnected layers of artificial neurons, where each neuron learns a component of a representation and successive layers combine these components to form increasingly abstract and useful representations of the data. Due to their high model complexity and large number of parameters, neural networks typically require substantially larger datasets and computational resources than traditional machine learning methods (Goodfellow, Bengio, and Courville 2016; LeCun, Bengio, and G. Hinton 2015).

Deep learning (DL) is a subfield of machine learning that focuses on a specific class of neural networks characterised by multiple layers. LeCun, Bengio, and G. Hinton (2015) describe deep learning as the study of “[...] computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction.” Accordingly, deep learning can be understood as the study of highly complex neural networks with many layers that learn hierarchical representations of data.

The specific arrangement of neurons and layers in such networks is described as the neural network architecture, which determines how information flows through the model and how representations are learned. There exist numerous such architectures which perform well at different tasks. Table 2.1 provides a non-extensive overview of relevant architectures with additional context.

Architecture	Domain	Source
Feedforward Neural Networks	General machine learning	(Rumelhart, G. E. Hinton, and Williams 1986)
Convolutional Neural Networks	Computer vision	(LeCun, Bottou, et al. 1998)
Recurrent Neural Networks	Sequential data	(Elman 1990)
Long Short-Term Memory Networks	Long sequence modelling	(Hochreiter and Schmidhuber 1997)
Generative Adversarial Networks	Generative modelling	(Goodfellow, Pouget-Abadie, et al. 2014)
Transformer Networks	Natural language processing	(Vaswani et al. 2017)
Graph Neural Networks	Graph-structured data	(Kipf and Welling 2017)
Diffusion Models	Generative AI	(Ho, Jain, and Abbeel 2020)

Table 2.1. A non-exhaustive overview of neural network architectures and their primary application domains.

2.2 Emergence of Generative AI

Language modelling was commonly performed using n -gram models, which predict the next word based on a fixed-length sequence of preceding tokens (Halevy, Norvig, and Pereira 2009). This approach inherently limits the available context, leading to suboptimal performance, particularly for long-range dependencies. To address this, Bengio et al. (2003) introduced neural probabilistic language models. These models leverage distributed representations to generalise across similar contexts, effectively capturing an exponentially large number of semantically related sequences. These models can be viewed as probabilistic next token predictors. However, a key limitation identified was the computational complexity of training such models, as “training such large models (with millions of parameters) within a reasonable time is itself a significant challenge” (Bengio et al. 2003). Therefore, the model’s architecture inhibited its scale in terms of model size as measured by its parameters, as well as limited its performance in terms of generation quality.

The constraining architecture problem was solved by the introduction of the *transformer* architecture (Vaswani et al. 2017). Unlike earlier sequence models based on RNNs or LSTM networks, the transformer relies on a mechanism called *self-attention*, which allows each token in a sequence to attend to—or consider—all other tokens in the input *simultaneously*. This design enables efficient parallel computation and improves the modelling of long-range dependencies in text.

Building on these architectural advances, modern language models are constructed as large neural networks trained to predict the next token in a sequence (Radford, Narasimhan, et al. 2018). These models leverage transformer architectures (Vaswani et al. 2017) to learn complex statistical patterns in language, enabling them to generate coherent and contextually relevant text (Radford, Narasimhan, et al. 2018; Bengio et al. 2003).

While early transformer-based language models already demonstrated promising generative capabilities, their performance remained constrained by limitations in model size, training data, and computational resources. Kaplan et al. (2020) show that model performance follows scaling laws as a function of these three factors. Consequently, improvements in architecture alone were insufficient; substantial performance gains required simultaneous increases in model parameters, dataset size, and computational resources.

Once this was understood, rapid scaling across all three dimensions followed. Radford, Narasimhan, et al. (2018) introduced the Generative Pre-trained Transformer (GPT) in 2018 with 117 million parameters, followed by GPT-2 in 2019, which increased the model size to 1.5 billion parameters and was trained on approximately 45 GB of text data (Radford, Wu, et al. 2019). This increase in scale led to significant improvements in generative performance. The trend continued with GPT-3 in 2020, which further expanded the model to 175 billion parameters and approximately 570 GB of training data (Brown et al. 2020), enabling advanced capabilities such as reasoning, translation, and code generation.

2.3 State of Practice in Generative AI

The rise of generative AI has not gone unnoticed by industry, where its unique capabilities have begun to disrupt a wide range of sectors. The technology has been applied across several high-level use cases, including automation, content generation, decision support and analysis, personalisation, data generation and simulation, knowledge work augmentation, creativity, and autonomous systems (Patel et al. 2026). Fig. 2.1 presents generative AI applications supporting digital automation across five industry sectors.




Figure 2.1. "Generative AI Applications in industry for Digital Automation", adopted from (Patel et al. 2026)

However, the adoption of generative AI extends far beyond these domains. Recent research highlights its growing integration in areas such as software engineering (Felder et al. 2026; Looi 2026), supply chain management and logistics (Sharma, Kakkar, Hooda, et al. 2026), marketing (Salih et al. 2026), finance (Eisfeldt and Schubert 2024; Metzger, O'Reilly, and Mac an Bhaird 2025), manufacturing (Boonmee, Mangkalakeeree, and Jeong 2025; Sharma, Kakkar, Hooda, et al. 2026), and construction (Taiwo et al. 2025). Consequently, generative AI can be characterised as a general-purpose technology whose capabilities enable transformation across diverse sectoral contexts.

Its applications not only span numerous domains but also multiple functional domains within organisations, including operations, marketing, finance, human resources, information technology, customer service, and supply chain management. Fig. 2.1 already provided an adequate overview of such low-level use cases within five sectors. Each low-level use case that finds its value through the application of generative AI is considered relevant. These applications are implemented by applying generative AI in any form—whether by adapting a foundational model or accessing a SaaS application. To exemplify, Table 2.2a summarises key use cases of generative AI in digital marketing together with commonly used tools, adapted from (Salih et al. 2026). Similarly, Table 2.2b presents representative use cases in software engineering and the tools supporting them, adapted from (Felder et al. 2026).

Table 2.2. Examples of generative AI tools across different application domains.

(a) Digital Marketing (Salih et al. 2026)

Use Case	Tools
Content Creation	ChatGPT, Jasper.ai, Writesonic, Frase.io, Canva Magic Write
SEO Optimisation	Frase.io, Writesonic, ChatGPT
Social Media Content Creation	ChatGPT, Canva Magic Write, DALL·E 2, MidJourney, Adobe Firefly, Fotor
Image Generation	DALL·E 2, MidJourney, Adobe Firefly, DeepArt, Fotor, NVIDIA Canvas
Video Production	Runway ML, Synthesia, Pictory, Lumen5, DeepBrain AI
3D Brand Visualisation	NVIDIA Omniverse, NVIDIA Canvas
Content Repurposing	Pictory, Lumen5, ChatGPT
Personalised Marketing	ChatGPT, DeepBrain AI, Adobe Sensei
Campaign Automation	Adobe Sensei, Runway ML, Canva Magic Write

(b) Software Engineering (Felder et al. 2026)

Use Case	Tools
Development Support	ChatGPT, Gemini, Claude
Local Model Deployment	Ollama, Mistral
IDE Code Assistance	GitHub Copilot, Cursor
Enterprise AI Tools	Company-specific proprietary GenAI tools
Code Completion	GitHub Copilot, ChatGPT, Cursor
Learning and Training	ChatGPT, Gemini, Claude
Bug Fixing	GitHub Copilot, ChatGPT
Testing Assistance	GitHub Copilot, ChatGPT

2.4 Theoretical Foundations of Technology Adoption

Generative AI as an innovation must be adopted to potentially unlock its business value. To that end, technology adoption must be understood. Such research in broad terms is governed by different stages of the innovation-decision process—ranging from knowledge, persuasion, decision, implementation to confirmation (Rogers 2003). Concepts introduced to investigate various aspects of these stages include, but are not limited to, adoption, acceptance, appropriation, and diffusion.

Adoption research fits into the third step of the decision-making process. The adoption process is when *“an individual (or other decision-making unit) engages in activities that lead to a choice to adopt or reject [an] innovation”*, where adoption is *“a decision to make full use of an innovation as the best course of action available”* (Rogers 2003).

Research studies adoption through a handful of perspectives. First, the determinants of adoption are frequently discussed. This research pertains to factors influencing the decision to adopt a new technology. Second, barriers and enablers of adoption are frequently investigated as well. Any factor that makes adoption more difficult or easier is of interest for this area of research. The adoption rate and the overall extent of adoption is likewise a crucial pillar of adoption research. This relates to understanding the breadth and depth of technology’s reach within a social system, and how this changes over time.

The determinants of adoption research fundamentally seek to answer the “why one adopts?” question. It identifies determinants and their relationship to usage, or more often the behavioural intention to use; where the latter serves as a proxy to the former. This line of inquiry has given rise to two prominent streams within the literature (Venkatesh, Morris, et al. 2003). Fig. 2.2 provides an overview of the technology adoption models that will be discussed next.




Figure 2.2. Technology adoption determinant theories

2.4.1 Individual frameworks

The first stream focuses on individual level adoption, examining the determinants of actual usage where Behavioural Intention (Ajzen 1991) to use a technology is often employed as a proxy. These models primarily explore cognitive, attitudinal, and motivational factors influencing adoption decisions of individuals, specifically.

The theoretical beginning of this stream rests in the *Theory of Reasoned Action (TRA)*, which identified that voluntary behaviour is preceded by the Behavioural Intention to conduct it (Fishbein and Ajzen 1975). Two main drivers of this intention exist, namely the individual's attitude towards the behaviour and societal norms surrounding it. The assumption is that the individual is free to choose any course of action. With the introduction of the *Theory of Planned Behaviour (TPB)*, a third one was appended to the already existing two factors; perceived behavioural control. This third factor was introduced to account for variation in behavioural intention caused by the individual's control to perform the action in question. Control was proxied by perceived control (Ajzen 1991). Taken together, individuals act as they have intention to do so. This insight laid the necessary foundations to study technology adoption through the lens of behavioural intention and its determinants.

The *Unified Theory of Acceptance and Use of Technology (UTAUT)* posits that the variation in technology use can be explained by four core determinants and moderators (Venkatesh, Morris, et al. 2003). It has four main constructs: Performance Expectancy, Effort Expectancy, Social Influence, and Facilitating Conditions. The first three are associated with Behavioural Intention. Whereas both Behavioural Intention and Facilitating Conditions have an effect on Use Behaviour. Apart from direct effects, there are four moderators of these relationships. Age plays a role in all four relationships, gender only in the first three, experience in the last three and Voluntariness of Use moderates the relationship between Social Influence and Behavioural Intention.

The *Technology Acceptance Model (TAM)* explains that the use of technology from an individual perspective is to an extent explained by two antecedents, perceived ease of use and perceived usefulness (Davis 1989). The former of which is partially causally related to the latter; if a technology is seen as being easier to use, it will also be perceived as being more useful. With subsequent research resulted in refined versions; TAM 2.0 (Venkatesh and Davis 2000) and TAM 3.0 (Venkatesh and Bala 2008).

Another perspective is provided by the *Technology Readiness Index (TRI)*, which measures individuals' propensity to embrace and use new technologies (Parasuraman 2000). Rather than directly measuring usage, the model captures psychological readiness toward technology. The model identifies four underlying dimensions, grouped into two categories: drivers and inhibitors. Optimism and innovativeness function as drivers that positively predispose individuals toward technology adoption, whereas discomfort and insecurity act as inhibitors that restrain adoption. Since it is an index, after some time it is in need of calibration. This resulted in *TRI 2.0* (Parasuraman and Colby 2015). Whereas the initial instrument consisted of 36 items, TRI 2.0 was streamlined to 16 items, of which 11 were retained from the original scale. Although the core purpose of measuring individuals' technology readiness remained unchanged, the revision was motivated by the recognition that measurement scales remain valid only as long as the constructs they capture are relatively stable.

2.4.2 Organisational frameworks

The second stream emphasises organisational level adoption, analysing the structural, strategic, and environmental determinants that shape firms' decisions to adopt new technologies (Venkatesh, Morris, et al. 2003).

A highly regarded framework explaining how an organisation's context influences technology adoption and implementation was originally introduced by Tornatzky and Fleischer (1990), while Baker (2012) provides a more structured account of the framework. The *Technology–Organisation–Environment (TOE)* framework posits that technological innovation is shaped by three contextual dimensions: the technological context, the organisational context, and the environmental context. Factors within each of these dimensions may either facilitate or hinder technology adoption, depending on the specific circumstances of the firm and its external environment. A key strength of the TOE framework lies in its flexibility, as the constructs within each context can be adapted to suit the specific technology and setting under investigation.

DiMaggio and Powell (1983) introduce the *Institutional Theory*, which posits that organisations operate within an organisational field in which they are subject to institutional pressures over time. These pressures are categorised as coercive, mimetic, and normative. Coercive pressures arise from formal regulations, legal requirements, and dependency relationships. Mimetic pressures emerge under conditions of uncertainty, where organisations imitate perceived successful or legitimate peers. Normative pressures stem from professional standards, shared norms, and educational backgrounds. Collectively, these mechanisms lead to institutional isomorphism, increasing structural similarity across organisations within a field.

The *IS Adoption Model for SMEs* proposed by Thong (1999) addresses two fundamental questions in information systems (IS) adoption within small and medium-sized enterprises: (1) which variables influence the decision to adopt an IS, and (2) which variables determine the extent of adoption after the decision has been made. The study identifies four relevant contextual dimensions. First, CEO characteristics, including innovativeness and IS knowledge, play a central role. Second, IS characteristics—specifically relative advantage, compatibility, and complexity—significantly affect adoption. Third, organisational characteristics include firm size, employees' IS knowledge, and information intensity. Finally, the environmental context is represented by competitive pressure. It shows that these factors influence both the decision to adopt and the extent of adoption, albeit with varying strengths. The decision to adopt is most strongly influenced by firm size, while relative advantage, compatibility, and both CEO and employee IS knowledge also exhibit significant positive effects. Regarding the extent of adoption, firm size and employees' IS knowledge have strong positive effects, whereas information intensity has a moderate influence.

The *Innovation Diffusion Theory (IDT)* proposes the innovation-decision process, which has the following constituents: knowledge, persuasion, decision, implementation, confirmation. Fundamentally, the theory aims to explain the manner in which innovations are diffused in a social system. Relative Advantage, Complexity, Compatibility, Visibility and Traceability are proposed as determinants of the adoption decision. All adoption decisions accumulated are the adoption extent for a given innovation. Moreover, it introduces the S-shaped adoption curve, showing how individuals can be classified as innovators, early adopters, early majority, late majority and laggards, depending on the time of their adoption. The adoption rate was likewise mentioned, which is "[...] the relative speed with which an innovation is adopted by members of a social system." (Rogers 2003).

2.5 Generative AI adoption

Generative AI adoption research has provided insights on individual and organisational adoption dynamics. First, adoption is discussed from an individual perspective using UTAUT as the organising lens. This is because the framework integrates constructs from several earlier individual adoption models, making it a pragmatic choice not to introduce unnecessary complexity. Second, the organisational perspective is discussed using the TOE framework, complemented by institutional theory. This is because TOE's three constructs can be adapted to the specific technology and organisational setting under study, allowing a single framework to capture three broad, relevant factors. Institutional theory enhances this view by placing more importance on external forces shaping adoption, which may be diluted in TOE as it distributes attention across three dimensions.

2.5.1 Individual Perspectives

Performance Expectancy (PE) is the most robust individual level predictor in the generative AI adoption context, however not universally so. Numerous studies state that it is one of the most, if not the most, significant determinants in their respective contexts (Smutny and Sudzina 2025; Fang et al. 2026; Xia and Chen 2025; Biloš and Budimir 2024). In contrast, Kim, Blazquez, and Oh (2024) report that it is not significant in early stages of generative AI adoption. In other studies it was significant, yet not overwhelmingly so (Camilleri 2024; Xin, C. Zhang, and Peng 2025).

Two other crucial determinants of the technology's adoption, yet less important than PE, exist. First, Social Influence (SI) is highly relevant (Xin, C. Zhang, and Peng 2025; Smutny and Sudzina 2025; Kim, Blazquez, and Oh 2024; Fang et al. 2026). However, the studies where SI was found to be most significant are mainly from China. Therefore, it is likely to presume that in more collectivistic countries, SI is more important than in, for instance, western European nations (Fang et al. 2026). SI's influence may also depend on organisational rank of the employee (Xin, C. Zhang, and Peng 2025; Kim, Blazquez, and Oh 2024). Second, Effort Expectancy (EE) is of similar importance as SI. It is crucial in early adoption of less experienced users. Kim, Blazquez, and Oh (2024) finds that it is crucial in early adoption, yet it fails to be significant in experienced or habitual samples (Biloš and Budimir 2024; Fang et al. 2026).

Consequently, the influences of PE, SI and EE on BI are likely to be context-dependent. They vary depending on adoption stage, country, familiarity with generative AI, and organisational rank of an employee. The organisational context likely influences individual adoption determinants as well. This is because organisational hierarchy moderates SI (Xin, C. Zhang, and Peng 2025), supervisory support drives SI (Kim, Blazquez, and Oh 2024) and task-tool fit shapes PE (Xia and Chen 2025).

2.5.2 Organisational Perspectives

Relative Advantage (RA) of generative AI is the most consistent technological driver of organisational adoption (Hughes et al. 2026; Xu, Ramayah, and Shi 2025; He et al. 2026). Complexity, conversely, was found to inhibit adoption, a barrier that surfaces somewhere else in a domain-specific form as the immaturity and lack of specialisation of current tools (He et al. 2026). Compatibility was presented to be a significant factor (He et al. 2026), however, this has been contradicted by Xu, Ramayah, and Shi (2025). Perceived value and perceived costs were also identified as influences (He et al. 2026), yet the connection between efficiency gains and cost savings has not yet been established.

Top Management Support (TMS) was a recurring organisational driver (Xu, Ramayah, and Shi 2025; Tewari, Al Alawi, and Morande 2026; He et al. 2026), although it frequently operates as a moderator or mediator rather than a determinant in its own right. Xu, Ramayah, and Shi (2025) found that TMS mediates both the relationship between Organisational Readiness and adoption as well as the relationship between Competitive Pressure and adoption. A second, closely related factor is Internal Skills: Staff Skills (Hughes et al. 2026), Versatile Talent (He et al. 2026) and Technological Readiness (Tewari, Al Alawi, and Morande 2026) point to the same underlying constraint, which is the workforce capability needed to effectively utilise generative AI.

Organisational Change Capacity and Readiness capture internal adaptability more broadly. Hughes et al. (2026) reported that Change Capacity significantly moderates the effect of Complexity and Skills on adoption, while Xu, Ramayah, and Shi (2025) found Organisational Readiness to be a direct antecedent. Additionally, He et al. (2026) identified Firm Size, Culture and Infrastructure as organisational-inherent characteristics shaping adoption.

Competitive pressure was a consistent environmental driver (Xu, Ramayah, and Shi 2025; Tewari, Al Alawi, and Morande 2026). The regulatory environment, by contrast, is contested. Hughes et al. (2026) found regulatory concerns did not significantly affect adoption, interpreting this as a lag between environmental pressure and organisational response. Meanwhile, He et al. (2026) found copyright and regulatory ambiguity to be a strong brake on adoption. Market, Partner and Customer pressures were all noted as further environmental influences (He et al. 2026).

Beyond the partial coverage of the environmental context of the TOE framework, institutional theory offers a more direct lens on regulatory and external pressures. Institutional pressures were reported to positively influence adoption (L. Zhang et al. 2025). These pressures, as previously introduced, can be coercive, mimetic or normative (DiMaggio and Powell 1983). Regulatory uncertainty acted as a negative moderator, dampening the effect of these pressures on adoption (L. Zhang et al. 2025). This converges with the regulatory lag (Hughes et al. 2026) and copyright ambiguity (He et al. 2026). Consequently, regulatory uncertainty is said to inhibit adoption as viewed through both frameworks. Innovative culture, by contrast, was identified to positively moderate that relationship (L. Zhang et al. 2025). This aligns with the previously noted organisational culture construct. Notably, the studies underpinning this institutional and regulatory paragraph are all situated in the Chinese geographical context (Xu, Ramayah, and Shi 2025; He et al. 2026; L. Zhang et al. 2025), where a strong state and policy likely exhibit different pressures than in countries elsewhere. Hence, these findings may not generalise well to contexts outside of China where these pressures differ in strength or even direction.

2.6 Research Gap

Recent studies identify Performance Expectancy, Social Influence and Effort Expectancy as the core individual level determinants of generative AI adoption. However, their influence is strongly context-dependent, varying with adoption stage, national context, familiarity and organisational rank. At the organisational level, the TOE framework groups factors into technological (Relative Advantage as a driver, Complexity as a barrier, Compatibility as contested), organisational (Top Management Support, Internal Skills, Change Capacity and Organisational Readiness) and environmental (competitive pressure positive, regulatory uncertainty contested). Institutional theory adds that institutional pressures drive adoption while regulatory uncertainty dampens it. Crucially, these two levels seem to be intertwined: hierarchy moderates Performance and Effort Expectancy, supervisory support drives Social Influence, and task-tool fit shapes Performance Expectancy.

However, across the reviewed literature, individual and organisational adoption tend to be separately examined, even though the two appear closely linked. Despite that, the current literature rarely combines the individual, organisational and external factors, all of which are necessary to provide a comprehensive picture of generative AI adoption (Y. Yesuf and Z. Fields 2025). Several authors likewise specifically call for greater attention to the organisational contexts shaping individual adoption of the technology (Lim et al. 2025; Kumar et al. 2025). This is especially under-examined in the small to medium enterprise sector, which is relatively under-represented and faced with distinct factors and pressures. Therefore, it is unclear how organisational and individual barriers and enablers shape each other, and how generative AI is adopted in SMEs.

As resource-constrained and organisationally less mature firms, SMEs lack insight into how to navigate their resource limitations and organisational challenges to adequately adopt generative AI. Beyond that, the resource-constrained nature of SMEs might be a perfect ground for innovation. These companies might develop their own practices, which are unknown to larger audiences. Consequently, specifically SMEs but also larger enterprises might benefit from understanding how generative AI adoption is influenced across both levels in a less forgiving environment. From a theoretical standpoint, individual adoption studies are conducted using frameworks as UTAUT and organisational adoption through TOE or Institutional Theory, with scarce integration across both levels. Generative AI adoption in the SME context is relatively less covered by existing literature, and there is a need for more qualitative research that will provide more in-depth explorations of empirical adoption practices.

This research aims to understand how organisational and individual barriers and enablers influence adoption of generative AI in small to medium enterprises, and provide a descriptive account of which tools are presently adopted for what use-cases. This research aim motivates the research questions introduced near the end of Chapter 1, which this study sets out to answer.

Chapter 3

Methodology

3.1 Research Context

Research on how individuals or organisations adopt technology is a fixture in information systems research. This research addresses how generative AI is being adopted by employees of small to medium enterprises. The first relevant context is the technology. AI has been topical since multiple years already, where generative AI only emerged a handful of years ago. The broader AI context is relatively well-covered in literature, whereas the more fine-grained context of generative AI is in development. Breakthroughs are occurring frequently to such an extent that it keeps academics and business executives waiting for what is next. To keep literature up-to-date, on-going research into this direction is required. Consequently, further investigations of generative AI adoption are warranted.

Furthermore, the organisational scope of generative AI investigations is important. Different people adopt technology differently, the same applies to organisations. Large enterprises have established governance structures, access to plenty of resources and are attractive to significant talent bases. They are less constrained in their adoption of generative AI, which might not reflect ingenuity, resourcefulness or novelty. In contrast, small to medium enterprises do not possess all of the luxuries of an established enterprise. This may result in diverging adoption practices between individuals employed in smaller companies and their larger counterpart at larger firms, stemming from the different organisational conditions. Consequently, the small to medium enterprise context necessitates separate and adequate attention.

3.2 Research Design

Guided by the interpretative view, the approach, design and methods employed in this study have been selected to adequately fit the research questions. Generative AI adoption in this research is studied through perceptions of individuals who are employed in small to medium enterprises. Fundamentally, their views on and assigned value to their experiences with generative AI, as experienced within their organisational contexts, form the basis of this research.

With that fundamental basis, the qualitative approach with an exploratory design is the most suitable. The technology's adoption dynamic in the SME context is sparsely covered by existing literature, and the technology itself is rapidly developing. Therefore, an in-depth exploration of its adoption dynamic within small to medium enterprises is warranted through a qualitative approach with exploratory design.

3.3 Sampling

To obtain data convenience sampling was employed primarily through the author’s personal network and, secondarily, LinkedIn and cold e-mail outreach. This was warranted by the author’s limited access to potential candidates.

Limitations were imposed on who would and would not qualify to be a potential participant. To this end, the following inclusion criteria were defined:

- The participant was required to be either full-time or part-time employed in a small to medium enterprise, classified as having between 1 and 250 employees (Rijksdienst voor Ondernemend Nederland 2025), at the time of the interview.
- The participant was required to use generative AI for work-related tasks within the SME in which they were employed and the interview was about.
- The participant was required to be fluent in either Dutch or English.

3.4 Data Collection

Semi-structured interviews were selected as the data collection method. Interviews were chosen over alternative qualitative data collection methods such as focus groups because they capture each participant’s individual perspective without pressure to conform present in group settings. Furthermore, as participants were drawn from various organisational settings and dispersed across locations, individual interviews were the most pragmatic choice. Furthermore, semi-structured interviews were specifically chosen as they support both a priori formulated questions, which allow for consistent data collection across the sample and flexibility to deviate from the pre-defined structure often necessitated by the exploratory design.

The interviews were allowed to be conducted both in-person as well as through remote means—a Microsoft Teams meeting—and were allowed to be held in either Dutch or English. Prior to an interview, the participants were obligated to provide consent, either verbally or in writing, for recording the interview and the subsequent use of the recording for research purposes. A consent form was communicated, explaining what data will be processed, how it will be processed, where it will be stored and for what duration, among other topics. The consent form is attached in Appendix A.

The structured questions in the interview guide were formulated as only open questions, and divided across two major chains of thought. The first one was ‘Generative AI use at work’ which had questions surrounding specific tools and use cases. The second one was ‘Barriers and Enablers’. It had questions on the company’s perspective on generative AI; barriers and enablers; and social influence surrounding its adoption. The interview guide is located in Appendix B.

3.5 Data Analysis

Thematic analysis with an inductive approach (Woo, O’Boyle, and Spector 2017) was used to analyse the interview data. No existing codebook was made use of, and each newly coded interview contributed to the emerging codebook. The interviewing and data preparation for data analysis were interleaved. Upon the completion of an interview, it was soon after manually transcribed and put into a standardised format. The interviews were transcribed in their own language, however, the codes were standardised to be in English.

A modified version of the progressive coding process (Braun and Clarke 2006) was employed using Atlas.ti as the facilitating tool. This altered process is presented in Fig. 3.1. The coding process was conducted in batches of five interview transcripts at a time. The process commenced with a loop consisting of familiarisation and open coding ending when all batch transcripts had been processed. The former allowed for general impressions of the underlying interview data to be formed and the latter for standardisation of concepts across quotations spanning various interviews. Once the batch had been processed, the refined phase with renaming, merging and grouping of codes ensued. After the first two batches, patterns and themes began to emerge from the data. The patterns were cemented by using smart code groups which allow for structured queries combining different codes to be formed. This is elaborated upon in the next section. Themes started to form as well, capturing strong, repeated patterns across the interview transcripts. Finally, the remaining batches were processed and, again, refinement followed. At last, themes were investigated on a deeper level. All relevant codes were put into their respective theme and sanity-checked on whether they fit the theme based on their underlying quotations.




Figure 3.1. The modified inductive coding process

3.6 Coding Scheme

The inductive coding process produced the following three lenses for tools, use cases and factors, one lens each. It also resulted in five themes of which none were defined prior to the data collection process. First, the tools were found to have three attributes: architecture, generality and domain. Table 3.1a elaborates on these three attributes. Second, the use cases were abstracted into six super-types of use cases. Table 3.1b presents these and provides additional details. Third, each factor influencing adoption was categorised into its effect on adoption, the formality which governs it, and the level at which it operates. Table 3.1c covers these three at length.

Table 3.1. The inductively developed classification scheme: tools, use cases and factor lenses.**(a)** Tool classification

Attribute	Value	Description
Architecture	Stand-alone	Accessed through its own dedicated interface, such as ChatGPT.
	Embedded	Accessed within the interface of a separate host application, such as Copilot in Microsoft 365.
Generality	General-purpose	Supports many use cases across multiple domains, such as Claude.
	Specialised	Supports a single use case or domain, such as Claude Code for software development.
Domain	Conversational LLM	Chat-based assistants used via natural-language dialogue, such as ChatGPT.
	Productivity & collaboration	Assists office, document, and teamwork tasks, such as Copilot in Microsoft 365.
	Code & development	Supports software development tasks, such as Claude Code.
	Media & creative	Generates or edits visual or creative content, such as Adobe Firefly.
	Knowledge & transcription	Captures or organises spoken or written knowledge, such as Plaud.
	Business-function	Serves a specific business function such as CRM, finance, or HR, either as a dedicated tool like Jaicob.ai or built into an existing platform like Confluence RoVo.

(b) Use case super-types

Attribute	Value	Description
Super-types	Generate	Generating new content from a prompt.
	Transform	Converting, reshaping or improving existing content.
	Automate	Executing a repetitive task or workflow with minimal human involvement.
	Retrieve & analyse	Finding, extracting or interpreting information.
	Reason & deliberate	Supporting thinking, judgment or decision-making.
	Build	Constructing a functional artefact or system.

(c) Factor lenses

Attribute	Value	Description
Effect	Enabling	Supports, facilitates or encourages adoption.
	Blocking	Hinders, constrains or prevents adoption.
Formality	Formal	Embodied through official policy, procedure or provision.
	Informal	Emerges through unofficial norms, habits or social behaviour.
Level	Organisational	Governs or applies across the firm as a whole, or a substantial part of it.
	Individual	Pertains to a single person or their own behaviour.

Chapter 4

Results

4.1 Sample Overview

The interviews were conducted between 20 April and 3 June 2026 and amounted to 625 minutes of recorded material across the 19 interviews. The resulting sample is presented in Table 4.1. Case numbers are ordered chronologically by interview date. Fig. 4.1 shows that the sample consisted of micro (21.1%), small (26.3%) and medium (52.6%) firms, with 42.1% of their employees being technical and 57.9% non-technical.

Table 4.1. Participant Overview

#	Industry	Classification	Role	Country
1	Renewable energy	Small	Co-founder	Netherlands
2	Real estate	Micro	Head of Data	Netherlands
3	Museum / cultural	Medium	ICT coordinator	Netherlands
4	Retail / art supply	Micro	CEO	Netherlands
5	Information services	Small	AI R&D	Netherlands
6	Civil engineering	Small	Quality Assurance Manager	Netherlands
7	Staffing / recruiting	Medium	Content Marketer	Netherlands
8	Telecommunications	Medium	AI Automation Specialist	Netherlands
9	Data consultancy	Medium	Data Governance Consultant	Netherlands
10	Healthcare	Medium	Data Analyst	United Kingdom
11	E-commerce	Micro	Content Manager	Belgium
12	Financial services	Medium	Business Developer	Netherlands
13	Manufacturing	Medium	Credit Control Manager	United Kingdom
14	Healthcare	Small	Doctor	United Kingdom
15	Investments	Micro	Partner	United Kingdom
16	Retail / commercial	Medium	Customer Service Supervisor	United Kingdom
17	Professional services	Small	Head of Marketing	United Kingdom
18	IT consulting	Medium	Student Consultant	Netherlands
19	Information technology	Medium	DevOps Engineer	Netherlands

Figure 4.1. Sample distribution by SME classification and role type



4.2 Tools in Use

Tool adoption concentrated in a small number of general-purpose conversational assistants, with the remaining tools forming a long tail each adopted by only one or two interviewees. ChatGPT led with 18 adoptions, followed by Copilot (14) and Claude (12), while 35 of the 43 tools were each adopted in fewer than three interviews. Fig. 4.2 presents the most adopted tools with at least three adoptions across the sample, coloured by domain. Despite Claude already appearing in its general-purpose version, it reappears twice in its specialised form through its Code and Cowork variants. Conversational LLMs and Productivity & collaboration tools were 6 of the 8 most adopted tools, displaying a bias towards these domains. Knowledge & transcription and business-function tools appear only among less frequently adopted tools not shown here.




Figure 4.2. Most frequently adopted tools

Stand-alone, general-purpose tools were the most adopted category, with 49 adoptions across just 9 tools. The remaining categories had comparable adoption counts to each other, ranging from 19 to 22. General-purpose tools in both interfaces accounted for 68 adoptions across 13 tools, whereas specialised tools had 42 across 30. Fig. 4.3 presents the adoption counts and tool counts for each combination of architecture and generality attributes.




Figure 4.3. Tool adoptions across the four combinations of architecture and generality

The general-purpose tools were exclusively either conversational assistants or productivity and collaboration tools. In contrast, the specialised tools comprised a mixture of all six domains, with no single domain dominating. As previously mentioned, a small number of general-purpose conversational assistants led the adoption with a long tail of less frequently adopted tools. As Fig. 4.4 shows, the previously mentioned long tail predominantly consisted of a considerable number of specialised tools with low adoption. There also is a difference between the top two and the lower two sub-figures. The top ones display general-purpose tools of both the stand-alone and embedded kind whereas the lower ones show the specialised tools. The adoption count of the general-purpose tools is higher and concentrated in fewer tools than in specialised tools, as reflected by fewer but notably higher peaks of the top two figures in comparison with the lower two ones.




Figure 4.4. Tools by architecture and generality quadrant

4.3 Use Cases

Generative AI tools were primarily adopted to generate new content (19), retrieve & analyse information (18), and transform existing content (16). These tools were also used to automate (14) whole workflows or individual tasks and to support reasoning and deliberation (13) before or during decision-making. Less often, they were used to build (10) functional artefacts, predominantly computer programs, as shown in Fig. 4.5. Each super-type reflects the adoption count of that super-type’s underlying use cases. This will be elaborated upon hereafter.




Figure 4.5. Use case adoption by super-type

Distributed across the super-types, a total of 20 underlying use cases have been identified with 149 total applications. Despite the top 10 use cases being half of the total, they account for 112 applications (75.2%), whereas the remaining ones only for 37 (24.8%). Fig. 4.6 presents the top 10 most mentioned use cases across the sample, coloured by super-type. Drafting, research, advisory, retrieval were the most adopted ones. It is worth noting that no meaningful pattern emerged relating to their super-types or specific use cases.




Figure 4.6. Most frequently adopted use cases by number of interviews

4.4 Factors Shaping Adoption

The previous sections showed what tools are adopted and for what use cases. This section turns to what shapes whether and how these tools are adopted. First, the Origins and Diffusion of Adoption is discussed to understand how generative AI enters and moves through the organisation, along with its practical drivers. This lays the ground for the Human Factor of Adoption, which shows how adoption is concretely shaped by the individuals taking it up. Tool Provision and Its Workarounds then presents how organisations provision generative AI tools for employees, where the limits lie, and how those limits are circumvented. Following this, Rules that Help and Hinder discuss how companies constrain individual behaviour around the technology's adoption. Finally, Protecting Data and Judgement explains context around why the rules are imposed in the first place and how individuals voluntarily impose additional limits on their own use as well. Table 4.2 explains each theme by its definition. It likewise presents the number of underlying factors that comprise each theme and, for each, the amount of quotations that were coded with a factor. Throughout this section, each factor is annotated with the number of interviews it appeared in, shown as ($N = x/19$). When a factor acts in both directions, its enabling and blocking counts are given separately.

Table 4.2. Summary of the five themes shaping generative AI adoption in SMEs.

Theme	Definition	Factors	Quotations
Origins and Diffusion of Adoption	How generative AI enters and spreads through the organisation.	22	256
The Human Factor of Adoption	How individual disposition and skill shape adoption once the tool reaches the employee.	18	256
Tool Provision and Its Workarounds	How organisations provide tool access, its limits, and how individuals work around them.	23	186
Rules that Help and Hinder	How organisations choose to govern and bound generative AI use through rules.	16	222
Protecting Data and Judgement	How the risks impacting data input issues and output of generative AI are managed	10	181

Fig. 4.7 displays how the factors are distributed across lens attributes, by theme. The effect shifts from largely enabling to increasingly blocking; formality tends to be informal at first and become more balanced later; and where individual and organisational factors are initially not strongly pulled into one direction, they are overwhelmingly organisational in the last two themes.




Figure 4.7. Distribution of factors across the three lenses as a % of factors per attribute

4.4.1 Origins and Diffusion of Adoption

Adoption of generative AI starts somewhere. Consisting of 22 factors, this theme describes how the technology ends up in and how it is subsequently diffused through the organisation. The factors are predominantly enabling (92%) and informal (75%), and lean organisational (63%) over individual (37%), as presented in Fig. 4.7. This figure is not cited again, though each theme’s opening paragraph draws on it.

Generative AI mainly enters the organisations through the sourcing by employees ($N = 9/19$), rather than through formal procurement processes. In some cases, this is driven by individual experimentation that begins at home and carries over into work, as one interviewee noted *“It’s employees doing stuff at home or just generally being interested, [...]”* (#13). Others begin directly with work use, as a customer service supervisor put it, *“So I started using it myself anyway and then it sort of rolled out from there.”* (#16). Some organisations have highly technically skilled staff and put that capability to use by establishing a dedicated AI body ($N = 2/19$), which is partly tasked with finding new opportunities. At one firm, an AI Lab was set up last summer which frequently holds knowledge session where employees discuss their experiences. However, the lab is also funded to experiment with the technology. Hence, some novelty might enter through this way as well (#5). If a company does not have such internal capabilities, some obtain them externally. Consequently, external professional or consultancy support ($N = 4/19$) assists the technology in entering the organisation through external channels. This is generally achieved by sessions together, where one interviewee said, *“[...] every year we do bring in an external party to explain that kind of thing further and to highlight the possibilities, the opportunities that come with it.”* (#4).

Once the technology ends up in the organisation either through employees, dedicated labs or external means, it is diffused by various channels across the organisation. Spreading the technology primarily happens through informal conversations between employees ($N = 14/19$). These conversations occur all the time. A data analyst remarked, *“My equal, we chat about it all the time, we chat about what we can use it for, I give her tips and she gives me tips on things she is using it for”* (#10). The time at which these happen does not matter, as some occur during a lunch or a break (#13). Such dialogues are not limited to technical staff either. A business developer mentioned that non-technical employees also discuss how they use these tools, and that these conversations extend beyond verbal conversations to online channels, such as a dedicated Slack channel (#12). They also help surface opportunities that someone too occupied with their own work might otherwise miss. During work, people are often not actively thinking about the potential improvements generative AI could bring, and a colleague can change that. As one manager described:

Yeah because I sort of think that I’m head down and have to do what I have to do. Processing orders and all that, so that’s the sort of stuff I have to do. So, they make me see it in a different light some of the time. Have you tried this or tried that? Have you tried it for this reporting and stuff? So, sometimes when you’re focused on your work you don’t think about other potentials. So, definitely I use AI a lot more than just the scope I intended to, which is great.

— #13, Credit Control Manager, manufacturing, medium

Beyond conversations, employees are frequently supported by internal experts ($N = 12$). Sometimes a non-technical employee becomes the reference point for a tool or for best practices. To illustrate, a customer service supervisor became *“[...] the guy who understands it [...]”* (#16). He became the person colleagues turned to for what the tools could do. Technical experts play a similar role. A specialist routinely shows colleagues how to get tools like Claude or ChatGPT working, even going as far as to show how to build applications with Claude Code and Codex (#8). Where someone already is an expert, others might decide to become one. They learn a tool or a generative AI skill and become the team’s expert on it. A head of marketing described a colleague who chose to master Copilot’s best practices and then enabled the rest of

the team with that knowledge. (#17). Expertise does not have to be concentrated in one person, though. A firm may hold a mixture of experts, so an employee might consult several to learn about different topics:

I have sometimes sat down with my colleagues about how to do it, how do I write a good prompt. I now work with projects, it is something my manager has recently explained. What you can do is specify within a project that you want this and this and this, you clearly explain the constraints and subsequently ask your questions in the project. Those are things that I have learned from colleagues and my manager. Those really have been my drivers.

— #11, Content Manager, e-commerce, micro

Peers also observe one another ($N = 13$), which spreads adoption further and can be strong enough to convert initially sceptical employees. A head of data who had been against generative AI switched after watching his colleagues use it. After having done everything himself for two years, he eventually decided to try it out once he saw others doing so and the impact it had (#2). A particular form of this is internal showcase, where individual accomplishments are shared with others ($N = 7$). Some employees take pride in showing what they have built. A co-founder noted simply that people *“show the things they have built”* (#1). This extends beyond simple artefacts to fully functioning applications built by non-technical staff. One interviewee described a colleague, not particularly technical, who built an app that is highly used internally. In doing so, he showed others what is possible and that even non-technical employees can innovate in this way (#5). Hackathons were also mentioned as a channel ($N = 2$).

The effectiveness of the diffusion channels described above is shaped by further factors. First, the leadership’s disposition towards generative AI acts both as an enabler ($N = 16$) and a barrier ($N = 2$). Leadership does not need to use the technology to effectively support their employees, as a customer service supervisor remarked, *“[...] my line manager [...] does not use it that much but she knows that it’s useful [...] she asked me to use it [...] everyone in the chain sort of upwards from me is positive about it”* (#16). Hence, what matters most is visible endorsement rather than leaders’ own use. Permission and encouragement from the top are sufficient to support adoption. Again, this need not be direct encouragement and can occur through online means. For instance, a Slack channel in which a vice-president was messaging the team about adoption, encouraging employees to share their journey and discuss findings (#12). In contrast, leadership inhibited adoption in some cases. An AI automation specialist recalled that when AI was first introduced at the firm, parts of the management team took time to accept that it would help—in his own words, *“that does inhibit you”* (#8). This early-stage lack of support eventually shifted as management became more supportive over time.

Apart from having supportive management, an open and innovative culture ($N = 11$) affects adoption in a positive way. A culture of openness and encouragement gives firms an advantage. Where one interviewee remarked how readily use was permitted, *“The openness to using it, the encouragement to say that’s fine, get on board [...] If I asked someone [...] it was straight away yes”* (#16), while another noted, *“We are a very open company, if there’s something that we could be doing better that’s what we are happy to hear.”* (#13). A flat organisational structure ($N = 9$) is complementary to that, where having management close to the work floor has a positive effect. As one interviewee stated, *“what makes it easier is that we’re a very flat organisation [...] because we’re a flat organisation and the leadership is genuinely innovation-focused, your suggestions for certain approaches do get listened to”* (#5).

Moving outside of the organisation, external pressure also shapes adoption, specifically industry and market pressure ($N = 7$). Many employees are aware of their company's standing in its sector and have an intuitive sense of the broader market movements. A credit control manager noted that in manufacturing they must constantly stay ahead of the curve: as the competitors adopt AI, so must they (#13). A content manager reached the same conclusion, *"if at some point you don't go along with it and don't start using AI, you simply can't keep up"* (#11). This sense of a shifting market and industry shaped how employees view the future, captured by a head of marketing:

We are living in this world like similar to when the internet was created [...] but the people who will flourish the most will be the people who understand it, know how to work with it and know how to prompt it.

— #17, Head of Marketing, professional services, small

4.4.2 The Human Factor of Adoption

Where the previous theme focused on how generative AI moves through the organisation, this one covers the individuals doing the moving. It shows how employees feel about generative AI, how they cultivate skills and what enables or inhibits that process. It has 18 factors, which are mostly enabling (72%), and largely informal (89%). They lean slightly individual, with 67% individual and 33% organisational.

Employees' attitudes towards the technology have an effect on how it is adopted. The disposition of employees enables ($N=13$) more than it blocks ($N=9$). A content manager stated, *"you generally have two groups, people who like it and the people who hate it"* (#7). To which group one belongs is not clear cut. Some employees are more open to new developments in general, they immersed themselves in generative AI early-on and have been using it more than others since (#5). However, there is also fear surrounding artificial intelligence:

But also a fear of AI in general. It is unfamiliar, and that makes it unloved. On one hand you hear positive stories, and on the other the dark side of AI as well.

— #3, ICT coordinator, museum/cultural, medium

Apart from general openness, some employees are driven by something stronger. This is an intrinsic motivation to employ generative AI ($N = 15$). Among technical staff, it is no surprise that the motivation is inherent, as an interviewee with an AI-background mentioned (#2). Meanwhile, employees with little to no technical background also find it extremely interesting and fun to use (#1). Where an ICT coordinator introspectively stated that, *"I sometimes wonder to what extent you really need technical skills to be able to probe further. It's often about being interested, having interest. That gets you a long way, rather than knowing what is technically possible"* (#3). Therefore, any employee could inherently be enthusiastic about the technology, which influences adoption. Some workers are so enthusiastic that they encourage others to use it, where a data analyst is motivating his superiors to use it and helps his older colleagues out (#10).

The need to help older colleagues hints at a broader pattern, where there is a generational divide in generative AI adoption. It is primarily a blocker among older employees ($N = 5$), where a CEO observed that *"you do see a generational difference, older employees find it a bit more daunting to work with and don't fully trust what comes out of it [...] compared to the younger guard, who deal with it more flexibly"* (#4). An observation shared by a credit control manager (#13). However, the pattern is not universal. In one case, it was actually reversed. Younger management members had to constrain an older executive who was overly enthusiastic, believing everything could be replaced by AI (#3). Older employees, therefore, tend to be more hesitant, though this does not always hold in practice.

Beyond age and attitude, adoption also depends on skills employees must build. The first is identifying where generative AI can add value ($N = 9$). This relates to finding use cases (#11), where even the models themselves can be employed to find them (#12). But also finding relevant tools to support tasks (#10, #11). Complementary to being able to find relevant tools and identify promising use cases, prompting skills have also frequently been noted as important ($N = 9$):

How to use AI properly, so what prompts to use rather than asking AI to do this. It is good to ask how to do this or how could I do this using this and this and this to get a better result. To get the right ingredients so that the recipe works.

— #10, Data Analyst, healthcare, medium

Being able to identify use cases, find the right tools for the job and know how to effectively utilise them could be developed either through individual initiatives ($N = 9$) or organisational programs ($N = 9$). The former can be done by reading information sources such as blogs (#9, #2) or using LLMs themselves to explain concepts (#19), for instance. The latter is where organisations actively provide upskilling opportunities. This can be done with training sessions, however, these can be quite limited. One interviewee stated that, “[...] we were introduced to Copilot that was mandatory. [...] But other than that everything else is learn as you go sort of thing. It is very much applied common sense.” (#16). Others are more effective and cover a wider range of topics:

We have a kind of traineeship and you have to complete the AI-literacy training regardless. It covers what AI is, how you use it, which different tools there are. You also get introduced to a platform called FIST, a sort of security-type platform where you learn the dangers of AI. Recognise this, never do that, when working with AI. Along two broad lines, in very broad strokes: how do you use AI and what can you do with it, and what are the dangers of AI and what should you watch out for.

— #18, Student Consultant, IT consulting, medium

If a company does not provide internal training sessions, it could offer an upskilling budget ($N = 3$) for employees to attend external ones. One interviewee is going to such a training later this year, which covers how to use generative AI tools, and how to get the most out of them (#17).

While training sessions may develop generative AI skills, pre-existing digital skills also shape adoption to some degree by influencing how effective someone learns new technology-related topics. Employees who are less digital skilled struggle more with adoption ($N = 8$) than those who already possess those skills ($N = 6$). From one side, a customer service supervisor expressed that his hands-on colleagues in the workshop, who have little to no digital competencies, have trouble understanding it: as he puts it, “*They understand hammers and saws and things, but less things about computers. I think it is a generational thing and also technologically savvy people are much more embracing of it whereas the practical people less so*” (#16). Moreover, as an AI automation specialist recollects, the technical employees have enough skills to automate their own processes using generative AI, whereas the individuals with less experience have more difficulty (#8). From another side, existing competencies with digital tools could help adopt generative AI. A partner said that she and her colleague are strong in Excel and financial modelling. This makes them adequate at knowing how to evaluate AI in the financial domain (#15).

4.4.3 Tool Provision and Its Workarounds

Disposition and skills only translate into use if the tools are within reach, and that access is something the organisation partially controls. Hence, this theme presents how enterprises provide access to the tools, and how employees at times make more with less. It also discusses how firms decide to limit that

provision through different ways and how employees in response find ways to circumvent those restrictions. Consisting of 23 factors, most are organisational (70%) and informal (61%). There is a roughly equal mixture of enabling (52%) and blocking (48%) factors.

First, some enterprises choose to directly pay for all tools and provide access ($N = 18$). In one case the company was paying for Gemini and Claude (#12), whereas in another it was paying for three tools (#7). Sometimes, a firm might only pay for one tool. According to the employee, even paying for a single tool definitely lowers the barrier to adopt such tools (#6). He mentioned that otherwise he would have to pay large amounts of money for them. This does not only regard basic subscriptions, as some organisations offer the maximum ones. A head of data first stated, *“At least the basic version, and they built on top of that.”*, which was followed up by, *“But there is an option to take out a max subscription.”* (#2). Apart from direct provisioning, organisations decided to provide vouchers for tools. However, these do have to be on an approved tool list. To exemplify, an interviewee explained that, *“[...] at the company where I work, you have a 20-euro credit to take out an AI subscription, to experiment lavishly with it [...]”* (#9). A specialised version of tool provisioning is embedded tooling in existing applications, which makes the tools easier to access ($N = 12$) for some and harder for others ($N = 5$). A few organisations are already paying for either the Microsoft or the Google suite, and employees use their embedded tooling. A data analyst mentioned, *“I think we have the licence to use all the Copilot stuff”* (#10), whereas a CEO said, *“[...] we have Gemini, but for us that’s part of a Google Cloud package.”* (#1).

Whichever form provision takes, access is rarely uniform across the organisation. Firms allocate it by role to make the most of their resources, scaling subscriptions to who the employee is ($N = 5$ blocking, $N = 3$ enabling). From the blocking side, an ICT coordinator explained that not everyone can obtain a paid account. He said that they are willing to give at least one person in every team an account, but only if it will be used enough (#3). This is decided per team. While other organisations decide per employed hours. A student consultant elaborated, *“they have a rule because we work with quite a lot of students. And put very simply, students get the least of the least.”*, where later he added, *“And indeed those dedicated licences were given straight to the permanent staff, and the students have to make do with shared accounts.”* (#18). From the enabling side, a business developer was given access to the, according to him, a middle-grade subscription. If he would need more, an approval process was established for additional tokens (#12). So, people were given different grades of subscriptions, depending on their roles.

Beyond which subscription role is granted, provision can also be loosened by leaving tool choice flexible ($N = 6$). Even where a firm pays for one tool, employees are free to use others alongside it:

[...] when they brought in the Microsoft suite they didn’t insist we use that alone. [...] They are more than happy for people to still use ChatGPT, [...], there’s been no pressure to switch over and exclusively use Copilot. Which is nice. It allows that flexibility.

— #16, Customer Service Supervisor, Retail / commercial, medium

The opposite also occurs: at times a specific tool is mandated and others forbidden ($N = 3$). A marketing head was initially barred from ChatGPT because Copilot was the paid-for tool, and she recollected the switch as, *“[...] really struggled. I likened it to, I have been on iPhone all my life and then to go to a Samsung, it’s very different”* (#17). Mandates are not always a matter of policy, they are sometimes imposed by technical constraints. A doctor explained that the network in his clinic is a closed system, where there are strong firewalls blocking any form of disallowed, external access, and because of this he could only access Copilot (#14). Although he would prefer to use other tools and did so when working from home.

Whether access is shaped by policy, firewalls, or anything else, the ceiling on accessing tooling is ultimately the cost ($N = 4$). However, this is perceived differently by management and the employees. The former sees it as a budget issue: “[...] putting all 130 of our employees on that, I find a) unnecessary and b) it would really take a chunk out of our budget.” (#3). While the latter perceives the token costs as important, “Ultimately, once you’re out of tokens, you can’t keep going. So the money component is a clear red line, because at some point it stops.” (#12). The token ceiling is where cost is actually felt. Even if a company does provision a paid subscription, the volume of accessible tokens remains a limiting factor ($N = 3$). One employee said that there is a specific number of tokens per five-hour window in his subscription and if he uses them too quickly, it becomes waiting till the window is up (#12). Developers also run into such limitations (#8). If no subscriptions were given out and employees had to use free tools, these also came with their own limitations in terms of functionality ($N = 5$). Some LLMs have a token model where the free tiers have a specific token amount allocated. This allocation is frequently insufficient because you quickly reach the maximum amount, as said by a DevOps engineer (#19). Claude is an example of such a tool with such a model. By contrast, ChatGPT only downgrades the LLM. A partner explained her strategy to navigate both models as follows:

Claude you can use it for free up to a certain point and then the tokens run out or you have to pay for it. And then you have ChatGPT which is somewhat nearly unlimited. In what you can keep feeding through it, it might downgrade your model but you can keep doing it.

— #15, Partner, Investments, micro

Just as the partner stretched limited access by switching between tools, others stretched it by sharing a single company-funded subscription across several people ($N = 3$). The account was rarely not directly shared by exchanging login information, though. Colleagues gained access indirectly by asking a person with a paid subscription to run their prompts, where a customer service supervisor noted that, “Someone said to me oh you’ve got the paid version, you have fewer limits on stuff. Can you do this for me. But other than that, yeah I have done stuff at their request I suppose.” (#16). The more common direct sharing of accounts also occurred between either two employees (#6) or even an entire team (#18). Yet, it is far from a perfect solution as token limitations only get magnified: “You get the best answer from the best model, in this case that’s Opus 4.8. Use it once and the entire usage of the shared account is gone” (#18).

When even shared access fell short or no corporate account was in reach, employees turned to self-funded (personal) accounts for work ($N = 5$). Employees in one case first relied on their personal accounts, where the organisation followed with introducing funded accounts (#5). Some companies still did not choose to do that and, as an interviewee put it, “[...] I obviously pay for ChatGPT myself anyway [...] There is no accounting department with endless budgets and such” (#16). So, if the company still does not make room in the budget, personal accounts could be used. But if limited room is made in the form of shared accounts, some may still overcome the shared account’s token limitations by introducing their own account. A student consultant mentioned that the shared accounts do not block him because, in his own words, “[...] I do sometimes use my own personal one. Yeah, I’m not going to make a fuss about the how and what. So yeah, not necessarily entirely what it’s supposed to be.” (#18).

4.4.4 Rules that Help and Hinder

Beyond how tools are provisioned, organisations also shape use through explicit rules. This theme explains how generally firms are permissive, by default. However, it also shows that rules do exist and they cut both ways, where they both enable and block adoption. It also discusses how control mechanisms are employed to enforce those rules. The theme has 16 factors where all are organisational, and primarily formal (69%). They lightly lean enabling with 56%, where 44% of the factors are barriers.

That default permissiveness is itself the starting point in many firms, formal restrictions are simply absent ($N = 10$). As there are no formal policies, employees view it akin to there being no rules to restrict the use of generative AI (#11). As one interviewee noted, *“it’s a real cowboy world, all sorts of things get thrown in, yeah, all sorts of things get thrown in, sometimes entire email conversations get thrown in, select everything and paste.”* (#1). And even if rules are there, they may be relaxed if one is responsible for one’s output, where a partner elucidated, *“[...] the rule is relaxed as long as [...] don’t get caught generating bad materials. You are still responsible for your output is probably the mantra.”* (#15). Employees see this absence of rules as something positive. When prompted about what has made it easier to use generative at work, he responded with the fact that *“[...] I get carte blanche from my manager. I think that in many companies there’s some kind of policy for it. [...] Well look, as long as I don’t enter the PIN code, he’s fine with it. So that does make it very easy.”* (#11).

Not every organisation is comfortable with leaving use entirely unchecked. Some accept it only with some degree of oversight, even when no formal policy exists ($N = 4$). It appears as a pseudo-policy, where employees are required to say what tools are used in an informal chat. This results in companies having a grip on what is being used to some degree. Moreover, even without such expectations, some employees take their own initiative and disclose that they were supported by generative AI. A quality assurance manager explained his approach:

[...] then I sometimes add that this conversation is summarised by AI. It’s a bit, well, yeah, it’s not written anywhere that I have to say it, but I am transparent about AI doing it, though people notice it themselves anyway. But I’ve never heard anyone say anything about it. [...] but we don’t have any written policy for it.

— #6, Quality Assurance Manager, Civil engineering, small

What appears informally as voluntary disclose is, in some firms, formalised into policy ($N = 3$). These are not restrictions on use, but there exist mechanisms that allow the organisations to have a grip on what is used by whom (#17). Apart from disclosing which tool is used, employees have to tell whether AI has been part of producing output. This appears as, *“[...] it was said that anything you do create, sort of tag it to say that this was created with AI.”* (#16).

Where such expectations remain unwritten, clarity breaks down. The problem with unwritten rules is ambiguity, where each employee interprets them differently ($N = 7$). Many companies leave it up to the employees, stating that they have the intelligence to be responsible, and if something feels off they should not do it (#11). However, as for privacy and data, there frequently is no clear line in the rules. As one interviewee explained, *“With data and privacy it’s not a very clear red line, sometimes it has an effect and sometimes it doesn’t. It’s your own judgement call.”* (#12). This dynamic becomes more complex with embedded tools that are becoming more integrated (#5). The core issue with unwritten rules is that their effectiveness hinges on interpretation. To exemplify, even though a data analyst stated the unwritten rule being not to input patient data as common sense, when prompted further about why they have not imposed strict controls for this rule, he introspectively uttered:

You know what, we have never thought about it, and now that you say it we might think about it. I need to bring that up, that’s good. No, in my mind, I would never do that but I guess someone else’s mind wouldn’t think the same, but I think it needs to be said. So yeah that’s a good point.

— #10, Data Analyst, Healthcare, medium

One way companies close this interpretation gap is through alignment meetings. Such meetings were stated to be effective in keeping employees aligned about what is and is not acceptable (#5).

Where alignment talk only partially settles the data privacy risk, some firms go further and introduce more formal data-input rules (enabling $N = 11$, blocking $N = 9$). The fact that rules become more clear is enabling, whereas their nature may block particular queries to be carried out. First, guidelines were established in some cases (#17). An ICT coordinator elaborated, *“So initially we actually set up a simple guideline. We can’t stop you outright, and we don’t want to ban it outright. But use ChatGPT very deliberately, so never put in personal data, no company-sensitive information, and make sure that where possible you don’t share your data for training purposes. That was the first small step.”* (#3).

Most tools are also being employed in their ‘safe mode’, which refers to not allowing data to be used as training data for the underlying model (#6). Furthermore, the fact that no person identifiable or any business critical data may be inputted is a strong rule. It limits what context could be prompted with, however employees have found solutions. Generalisation is applied to the degree where the prompt is no longer traceable to specific individual. As mentioned in one interview, *“But the main thing we were told was don’t put customer details in it, keep it generalised. You can put scenarios in there but don’t put Smith at 1 Road Lane, it was always keep it generalised.”* (#16). Despite the rules, some tasks still require inputting sensitive data. In one case, snippets of proprietary code were being inputted as carrying out the task would otherwise not be feasible (#5).

Beyond targeted data-input rules, some firms formalise their stance into fully written policy, which proves to be both enabling ($N = 6$) and blocking ($N = 6$). Some organisations were risk averse at first, starting with a restrictive policy and then loosening it with time. For instance, the policy was first restricting generative AI use to the extent where only management was allowed to use it. Then, as risks were better understood, the policy got expanded and allowed all employees to use such tools, yet still excluding trainees (#13). Policy may still be restrictive to the frustration of employees. An innovative interviewee expressed that, *“So I notice that it does sometimes cause me some frustration. That not everything is possible yet that I would actually like to be able to use.”* (#5). This signals that using generative AI for some things is directly prohibited by such policies, which is also supported by other organisations (#17). The case with policy is that it may have an adverse effect on adoption, as it may focus too much on restrictions and neglect how to enable employees to actually use the technology:

Yes, there is a policy, but it’s mainly restrictive, it mostly describes what isn’t allowed. But not much has been set up yet to help employees use GenAI in their daily work. That’s actually a shortcoming, I have to say. I think they’re falling behind on it.

— #19, DevOps Engineer, Information technology, medium

A policy on paper, however, is only as strong as its enforcement mechanism. Even when policy exists, organisations are selective about how strictly they enforce and how ($N = 10$). The overarching mechanism is restricting access to tooling. First, allow lists were used to force the use of tools that were perceived as reaching the mark with privacy, among other factors. For instance, Mistral’s Le Chat was chosen in one case as the designated tool over ChatGPT (#3). Even if tools are generally accessible, the initial download on a company’s laptop can be restricted through mobile device management (#19). In such a case, an employee needs to reach out to IT personnel to gain approval. Yet, an open gap exists, as most tooling is accessible through the web. In response, rarely but notably, closed systems are employed. In cases where such a system is already established prior to introduction of generative AI, organisations may use it to restrict access on the network level to web applications (#14).

Where access is restricted, a formal route to exceptions follows, which differs per company ($N = 7$). First, some companies do it through the approval of management. If an employee finds a potential tool, they will inform their manager. Then, it is decided from there. If the manager sees it as promising, a business case is written and sent up the management chain (#16). Approval may also be given as high as from C-suite executives (#8) or a designated AI governance body (#9). Where management considers the business case, IT considers it from a security perspective (#16). The fact that a tool is already known plays a role in their decision, as a student consultant stated, *“Yes, it also depends on how well-known the tool already is. Because let’s say, n8n was a fairly well-known name in the agentic AI space. [...] So introducing it was easier, and indeed the first client use cases came through soon after.”* (#18).

What these routes share in common is that the decision rarely rests with the user. Authority over generative AI tools is typically centralised ($N = 8$). General (technology) management is often responsible for making such choices (#12). The reason being, take a VP engineering for example, he has the necessary knowledge to decide whether a tool is safe or not. A second option is for the entire management team to decide, which in practice appears as the CEO, CTO and COO coming together (#8). Moreover, employees with budgetary discretion may choose by themselves (#6) and even for others (#17). An AI governance body can also be tasked with making such choices (#9).

Whoever holds the authority, decision is rarely made on internal grounds alone. External constraints—ISO certifications or a firm’s own commitments—shape and constrain what can be approved ($N = 5$). When organisations cooperate with different parties and have, for example, ISO standard requirements, generative AI becomes more restricted, as noted by an AI researcher, *“Well, the biggest point, I think, is that we work with a lot of sensitive information from housing associations, and because of that we have to comply with many ISO standards, which makes us more cautious by nature.”* (#5). He followed this up with, *“Well, I think mainly compliance, but harder is relative. It hasn’t become harder to request tools, but they do make it difficult.”* (#5). Compliance is in some cases less hands-on, where they may only give a lecture but not follow it up with anything more and leave the employees to their own devices (#15). Beyond compliance, data protection officers influence tool approval policy (#9), and frequently discuss risks as well (#19). In one company, even a dedicated AI compliance officer was employed (#18).

4.4.5 Protecting Data and Judgement

The rules, policies and controls described above all exist for a reason. To manage the risks organisations perceive in generative AI. This theme sets out what those risks are, how human-in-the-loop is the go-to solution and what employees consider as risks for them. Across its 10 factors, most are blocking (90%) and organisational (80%), with evenly distributed formal and informal factors.

Organisations primarily are concerned with data privacy risk ($N = 15$). The fact that they have data of individuals, either personally identifiable information or financial data, makes organisations more conscious about their generative AI utilisation (#5). For example, a business developer and his technical colleague were evaluating OpenClaw when they came to the following conclusion:

It’s all very nice, but he said it’s too dangerous for us because it can do so much on your computer, it basically takes over your whole system, so we’d rather you didn’t use it because we’re afraid it could leak all kinds of data.

— #12, Business Developer, Financial services, medium

This translation of privacy fear into a tool restriction was not unique to one firm. An ICT coordinator explained that despite repeated requests for Copilot, he followed SURFmarket advices and declined to allow it, restricting the use of American models because to the extent to where privacy is being respected

remains unclear (#3). Underlying this is a specific concern, which is enterprise data exposure risk ($N = 6$). It is the unintended side-consequence of granting a model read access to internal systems. Where enterprise systems are not neatly organised with proper separation, giving a model read access to the entire data store can cause accidental data egress: *“And things like copilot, they are not going to happen. And also because our SharePoint is not arranged in such a way that I can confidently say that everything is neat.”* (#3).

The worry is not only about the data itself but also about where it would land. It is the reputational or client risk it carries ($N = 5$). Companies hold themselves to high data privacy standards partly because it would cause reputational risk otherwise: *“[...] above all a great deal of concern for the data we use, because our clients would, well, really be put off if something were to leak.”* (#5). A head of marketing explained a different side of this risk. Their social channels rely on raw and fantastic content for their loyal customers. This makes them, as she claimed, the best performing version of their consultancy in the country. The worry is that their fans would be able to tell straight away that AI is being used (#17). This also applies to their business clients, as she stated, *“[...] we are always seen as the people who are the problem solvers, but if people think that AI is their problem solver it is potentially damaging for the company.”* The client-facing risk can translate into concrete restrictions:

[...] my colleague was a huge champion of using Gamma. But he came from a startup and venture capital background whereby things flow fast. The other two of us gave him feedback that it is just not good enough for institutional investor presentations and we stopped that.

— #15, Partner, Investments, micro

From the employee perspective, the reliability of output and its trustworthiness is a recurring concern ($N = 12$), in two distinct senses. First, whether the output can be trusted as correct. Output is rarely taken at face value. A head of data explained that figures are always treated with caution and verified rather than blindly trusted: *“[...] the figures that come out of it are always taken with a grain of salt. So everything is checked and not blindly trusted.”* (#2). A second worry is that AI output is recognisable as artificial. Some felt it “still looks too AI” to use for something like emails, preferring to write themselves (#10). In creative work, this concern runs deeper. The fear becomes losing the organisation’s own voice:

Because where our fears lie is that if we really start developing programmes for the public and use AI to do so, everything ends up homogenised. [...] the writers we employ, or at least the people who work on our texts, they are really anti ChatGPT or Mistral. Where does our identity go? So that is something that really lives here.

— #3, ICT coordinator, Museum / cultural, medium

Beyond those worries, they likewise fear deskilling and over-relying on generative AI ($N = 5$). In one case where individuals started relying on it a bit too much, to such a degree that *“it has become a bit of a joke in the office like I have written an email, no Copilot has written that email.”* (#17). Not everyone is comfortable with that reliance. Even though a content manager expressed that he is not a fan of these tools because they make him lazy, he uses them anyway whenever he does not feel like doing the work himself (#11). This concern is shared by management, a co-founder said, *“But whether, as a company, I’m happy that people no longer think for themselves and just uncritically copy everything, somewhat less.”* (#1).

Where deskilling concerns thinking, a related fear arises once the tools are allowed to act on their own. Employees frequently have reservations about the loss of control around AI’s actions ($N = 4$). An interviewee who was using Jaicob.ai allowed it to fish candidates across social media (#7). Yet, unknowingly,

the system sent 3,500 LinkedIn invites. This resulted in her account being temporarily blocked from sending invites. She felt the adverse effect and put in stronger guardrails. Another example is the head of data, who mentioned, *“If you accidentally send six thousand emails saying ‘test’ to some people, that’s a red flag.”* (#2). Even technical employees could easily make a mistake with autonomous systems, which could accidentally lead to removing repositories or egressing proprietary code (#19).

To manage these risks—from organisational concerns around data and reputation to unreliable output and unsupervised action—the most common safeguard is to keep the human in the loop for verification ($N = 15$). There initially were problems with verification, as it was not being conducted. An employee recollected that there was an increasing number of reports being delivered purely written by AI, where at this point they had to draw this use back (#3). With proper verification, there is always a check before an LLM’s output gets used. A head of data states that, *“if we draft an email based on figures, the email is always written and reviewed by a person first, and then a check is run over it to see whether it makes sense.”* (#2). However, it is important to possess the knowledge to actually verify the output’s correctness. As one employee stated, the models are not 100% accurate. Therefore, it is important that, *“And to have that knowledge yourself, so that you can take out the mistakes AI makes, because it does happen sometimes.”* (#8).

Beyond verifying AI’s output, some firms keep certain tasks in human hands entirely ($N = 6$). Some managers ask for direct human input, where the employee is expected not to use AI at all (#12). Prohibitive policies may also be drafted for specific tasks, as a head of marketing mentioned, *“There are some things that we see that are against company policy, so we won’t use it to write our social media posts. They have to come from us, they have to be organic”* (#17). This aligns well with the idea of personal service from smaller companies that clients enjoy, which was mentioned by an interviewee. He said that their unique selling proposition is personal service. He commented that, *“I work for a very small company, and one of our big strengths, our USPs really, is that we exist as a company built on personal contact. So yeah, then you quickly become wary of AI.”* (#11).

Limits are not only put on what AI can be used for, but also about which tools are permitted in the first place ($N = 8$). There is a sizeable push for European Sovereignty, where the models should either run in Europe or have their inference work done there. An interviewee described it as:

With us, everything has to be properly vetted first. It has to run in Europe, the inference really has to happen in Europe. So quite a lot depends on that, and because of it there have sometimes been tools I couldn’t use.

— #5, AI R&D, Information services, small

Despite many other organisations feeling an inclination towards European models, they frequently choose American ones because of performance. An AI automation specialist described that they had considered switching, yet the capacity of European models is not sufficient (#8). The same conclusion was reached at another firm, where an employee put it as: *“We’re leaning towards it more and more, it was a factor in the tool selection, only I think it was ultimately decided that it’s a lower priority, because the tools in Europe are just a bit less developed. But slowly, once there’s a tool that runs just as well as the current ones, the switch to Europe will happen.”* (#18) Alongside geography, the ethical positioning of the model provider also weighs on tool choice ($N = 6$). Companies manage risk, which include data privacy and reputational risks. So, the ethical positioning of the model provider may reflect on them (#5). But again, despite seriously discussing it, action does not follow.

4.5 Variation across Organisational Size

This section presents the tendencies that emerge when the data is partitioned by organisational size. It is important to note that the resulting bands are small and uneven, with medium firms (10) outweighing micro (4) and small (5) ones. Therefore, these findings should be regarded as illustrative rather than generalisable.

Micro companies tend to predominantly rely on stand-alone tools, while the use of embedded tools increases in small to medium ones. Embedded tools constitute only 11% of all tools that micro enterprises use. In contrast, of all tools small and medium firms use, 48% and 38% are embedded, respectively. Fig. 4.8 visualises the distributions of embedded and stand-alone tools, per each classification band. It reflects that small and medium individually do not overwhelmingly favour one architecture over the other, unlike micro SMEs. Such micro enterprises predominantly choose stand-alone tools, as indicated by the significant difference between the orange and green bars on the left side.




Figure 4.8. Share of embedded and stand-alone tool adoptions by organisational size

Second, factors shift from informal to formal as organisational size increases. Micro firms were strongly informal (70% informal and 30% formal), whereas small and medium firms were more balanced at roughly 40-60 and 41-59, respectively. The shift is most visible between micro enterprises and the larger ones, as presented by Fig. 4.9. It shows that informal factors outnumber the formal ones. But not evenly so as partitioned by organisational size. As size increases, the gap decreases. This indicates that larger firms steadily have more formal factors, yet still not more than informal ones.




Figure 4.9. Share of formal and informal factors by organisational size

The blocking factors transform from informal, individual to formal, organisational as size increases. In micro enterprises, 47% of the factors were informal and organisational, whereas 29% and 24% were informal, individual and formal, organisational factors, respectively. Factors in small firms were more balanced with 43%, 16% and 41%. This already reflects the shift of informal, individual to formal, organisational factors between micro and small firms. In medium companies, the distribution was 39%, 17% and 44%, showing that formal, organisational factors were the largest group. Therefore, micro firms are mostly informal, organisational factors, while in small firms it is roughly even and medium firms lean formal, organisational. Fig. 4.10 illustrates this progressive tendency.




Figure 4.10. Composition of factors by organisational size

The share of both enabling and blocking factors did not meaningfully vary across organisational size, as shown in Fig. 4.11. It presents the percentages of both enabling and blocking factors per organisational size. There are more enabling factors across all sizes. However, medium firms have slightly more blocking factors than micro and small ones as indicated by the decreases gap between enabling and blocking factors for medium firms on the right side of the figure.




Figure 4.11. Share of blocking and enabling factors by organisational size

Chapter 5

Discussion

5.1 Summary of Findings

This study explored what organisational and individual factors influence generative AI adoption in small to medium enterprises. Prior work has tended to treat individual and organisational adoption separately, despite indications of the contrary. This study instead examines how the two interact, offering a perspective that spans both levels. Drawing on 19 interviews across a range of Western European SME employees, the inductive analysis produced 89 factors grouped into five themes: how adoption originates and diffuses; the human disposition and skill behind it; how tools are provisioned and worked around; the rules that govern use; and how data and judgement are protected.

Generative AI predominantly enters organisations through employees experimenting on their own, but also somewhat through dedicated AI labs and external (consultancy) channels. It is diffused through informal conversations, internal expert support, peer observations, and internal showcase. Leadership support is associated with adoption, where positive endorsement is sufficient and leading by example through using such tools is not required. Interviewees attributed an open and innovative culture and a flat organisational structure as an enabling environment, where they described efficiency gains as a pull to adopt such tools while they appear to simultaneously perceive the tools as intuitive. Externally, competitive and market pressure, and regulatory concerns seem to influence organisation and employee adoption.

Tool adoption concentrates on a few general-purpose conversational (ChatGPT, Claude and Gemini) and productivity assistants (Copilot). A long tail of specialised tools exists, where each tool is adopted in only one or two firms. Such general-purpose tools were adopted 68 times, while specialised tools only 42 times. Meanwhile, the former consists of 13 tools, the latter of 30. Of the general-purpose tools, the stand-alone interface dominates adoption with 49 adoptions in contrast to the embedded one with only 19. Micro firms adopt for the most part solely tools with a stand-alone interface (89%), whereas small (52%) and medium (62%) firms opt for both interfaces in a more balanced way.

Firms provision resources unevenly and react after the point of entry. By employees, provisioning was found to lower the barrier to adoption. Under-provision drives workarounds, including tool switching, account-sharing and self-funded personal account use for work purposes. Governance is permissive by default. Formal policy is the exception and unwritten norms the rule. When formal structure does appear, it is double-edged. Some rules enable one role but constrain another. This appears primarily as token or access constraints. Policy is a reaction to risk, whether data privacy, reputational, output reliability or loss of control over autonomous generative AI systems. The dominant safeguard is human-in-the-loop verification and common sense.

5.2 Interpretation

Bottom-up adoption and the organisation's reactive role

Across the aforementioned themes, one pattern dominated. Generative AI enters SMEs through employees experimenting on their own and spreads through informal conversations, peer observation, and internal experts rather than formal roll-out. This aligns with Kim, Blazquez, and Oh (2024), which states that social influence plays a crucial role in adoption. The mechanisms described above are practical channels through which that influence operates. Informal conversations, peer observation, and internal experts are how colleagues come to perceive that those around them value and expect engagement with the technology. Hence, as literature established *that* social influence matters, this study shows *how* it operates. Adoption travels laterally and across the ranks rather than being pushed down the hierarchy. This positions the organisation as the setting in which peer adoption occurs rather than its initiator, and where social influence is a highly relevant factor when adoption is not mandated.

The relevance of social influence remains significant across the collectivistic and Western context (Fang et al. 2026). This study supports that finding, as it is rooted in the Western European context where social influence also appears. Therefore, it confirms what previous literature found in a similar context. Further, Xin, C. Zhang, and Peng (2025) finds that social influence is rank-dependent, which this study has not empirically observed. Managerial and non-managerial employees both mentioned similar experiences where their colleagues slightly influenced them in relation to generative AI adoption. A potential explanation is that the study is rooted in the Chinese context, where rank may carry more weight than in the Western European setting of this study. This suggests that bottom-up, peer-driven adoption may itself be context-dependent. It may be peer-to-peer in flatter, lower power-distance settings but more hierarchically influenced where rank strongly structures social influence.

Beyond social influence affecting adoption, this study finds leadership's support as a key factor. This conforms to present literature, where leadership support predicts adoption (Tewari, Al Alawi, and Morande 2026) and top management support mediates between competitive pressure and adoption as well as between organisational readiness and adoption (Xu, Ramayah, and Shi 2025). Where these studies established the effect, this study shows that what matters most is visible endorsement and clear permission rather than leaders' own use of the tools, suggesting it is the signal of legitimacy that positively influences adoption. Moreover, in settings where social influence is rank-dependent, the effect of leadership endorsement may carry particular weight for lower-ranked employees, for whom a signal from above legitimises adoption more strongly than peer influence alone.

As leadership support impacts adoption, the efficiency gains that generative AI provides also have been cited as a reason for adopting the technology. Managerial interviewees consistently welcomed the improvements it brought, particularly to less digitalised parts of the business. Both Xu, Ramayah, and Shi (2025) and Hughes et al. (2026) identify Relative Advantage as a driver of organisational adoption, which aligns with the present findings. Employees also were positive about their efficiency gain, mentioning them as a relevant reason to adopt generative AI. This is captured by performance expectancy identified in present literature (Smutny and Sudzina 2025; Xia and Chen 2025; Biloš and Budimir 2024). However, in settings of bottom-up, peer-driven adoption with no mandated tools, organisational Relative Advantage may be secondary to individual Performance Expectancy. Where an individual chooses tools based on the performance they expect, and the organisational relative advantage is realised only once Performance Expectancy in the eyes of the adopter is high enough. The technology's complexity, however, was not experienced as constraining since the interviewees found the tools intuitive. This may reflect how intuitive current tools have become, suggesting perceived complexity is lower than earlier framings imply.

Beyond internal factors, the external environment affects adoption through competitive pressure. This corresponds with previous findings (Tewari, Al Alawi, and Morande 2026). This study extends the context in which such pressure has been mentioned, as the previous study was conducted in a BRICS+ setting, signalling that competitive pressure may be relevant beyond contexts previously studied. However, where regulatory concerns did not significantly affect adoption in one study (Hughes et al. 2026) but did in another (L. Zhang et al. 2025), this study finds that regulatory concerns are applicable factors. Privacy and AI regulations were frequently noted during the interviews, and a few companies had employees—DPOs, compliance officers or even a VP engineering—tasked with interpreting regulations and forming unwritten rules or formal policy to shape generative AI use. Where Hughes et al. (2026) found a lag between regulations and practice adoption, this study indicates regulation may now actively shape practice, suggesting that the gap has begun to close.

Tool concentration and the generality advantage

Generative AI adoption concentrates on a few general-purpose assistants—such as ChatGPT, Claude or Copilot—followed by a long tail of specialised tools each taken up by a few firms. In digital marketing the same patterns emerged, where only 1 general-purpose tool was adopted and 17 specialised tools (Salih et al. 2026). In contrast, 4 tools of both types were employed in the software engineering domain (Felder et al. 2026), which differs from both the marketing and this study’s pattern. However, an explanation for this may exist. The previous broadly aligns with task-tool fit being the strongest predictor of adoption intent (Xia and Chen 2025). Where general-purpose tools fit software engineering’s tasks well, they are adopted over specialised alternatives that may not provide an advantage. Conversely, in digital marketing where the fit of general-purpose tools may be weaker, specialised tools prevail. This implies a default-and-fallback logic in tool choice. A general-purpose tool is the first resort, and specialised tools are reached for only when it proves insufficiently performant for a given task.

Moreover, the dominance of a small set of general-purpose tools is also consistent with performance expectancy driving habitual use (Camilleri 2024). Once a capable general-purpose tool meets the performance threshold, users entrench and stop seeking specialised alternatives. The long tail of specialised tools also might reflect their relative immaturity (He et al. 2026), where specialised tools must be mature enough to outperform an already-entrenched general-purpose tool’s default; a bar that is not trivial to clear. As general-purpose tools appear to have matched common tasks early and well, their habitual use may have become strong enough, and their task-tool fit sufficiently attractive, to entrench perceptions of their relative advantage over specialised alternatives. This hints at a first-mover advantage, where an early tool’s capabilities are perceived as good enough, it becomes difficult to dislodge. Dislodging remains possible, however, only when an early mover falls far enough behind competing alternatives in terms of performance.

The double edge of formal structure and risk-driven governance

Organisations tend to contribute relatively little at the points of entry and respond mainly afterwards. They provision access unevenly—some funding a tool while others several—and govern through rules that are permissive by default, where unwritten norms are the rule and formal policy the exception. While Kim, Blazquez, and Oh (2024) concludes that facilitating conditions have no significant impact on early-stage adoption, in this study they appear to be relevant. The provisioning of tools was found to affect adoption. Some interviewees stated that they would not be prepared to pay themselves for a subscription utilised for work purposes, while others directly highlighted that their company’s provisioning definitely lowered the bar to adopt generative AI tools. There were some who already applied these tools in their personal lives, and when not prohibited by the organisations, they did use it for work purposes. This signals that facilitating conditions may play a larger role in later-stage adoption, and especially tool provisioning seems to have a strong effect.

Where formal structure does appear, it is double-edged. The same mechanism enables one employee while constraining another. When policy becomes perceived as unnecessarily restrictive, employees may choose to disregard it. While this pattern is not directly addressed in the reviewed literature, the closest match is Xin, C. Zhang, and Peng (2025), which finds that standardised implementation lands unevenly across ranks. This study extends that observation. Adoption can diverge by role as a direct result of policy choices, where a policy favours some roles over others. Technical roles, for instance, frequently had access to greater token allowances—through direct purchase or higher-tier subscriptions—than non-technical ones. This produces a divergence in adoption dynamics introduced by policy itself.

Meanwhile, such policy arrives as a reaction to perceived risk—data privacy, reputational, output reliability and loss of control—managed above all by keeping a human in the loop to verify prompts, outputs and vibe-coded programs. This corresponds to existing research, which identified hallucinations and misinformation (Smutny and Sudzina 2025), source trustworthiness (Camilleri 2024) and privacy concerns (Lim et al. 2025) as factors affecting adoption. It implies that these factors are broadly shared across different contexts and shape adoption beyond the contexts studied before. However, loss of control has not been identified in examined adoption literature. Autonomous capabilities were probably less accessible or not performant enough when those studies were conducted, suggesting new risk factors may surface when new tool capabilities develop.

Tool selection was further shaped by a push for European digital sovereignty, where models were expected to run and perform inference within Europe. This reflects an institutional pressure of the kind described by DiMaggio and Powell (1983) and its generative AI application by L. Zhang et al. (2025). In practice, however, this pressure consistently lost to performance. Firms take European sovereignty into account within their tool selection decision-making process, yet still consistently opt for American models for their superior capabilities. This is a clear case of relative advantage overwriting institutional pressure, the sovereign preference persists as an intention but not as action until capability parity is reached.

5.3 Theoretical Contributions

This research advances how generative AI adoption is studied in three primary ways.

First, it helps close the existing research gap, in which adoption is separately studied at the individual and organisational level by providing a unified view spanning both levels. It does so by showing that the two levels are closely intertwined. In this setting, the risks organisations face constrained individual adoption through, for example, unwritten norms and, at times, formal policy. Moreover, organisational budgetary decisions enabled or constrained access to tools. This affects individual adoption by shifting facilitating conditions.

Second, it adds novel findings to adoption research situated in the small to medium enterprise context, which is a relatively under-studied area of adoption research. It presents that—in contrast to larger, resource-rich companies—small to medium enterprises and their employees find ways to make more with less, particularly around tool subscriptions. It also suggests that emerging differences arise across SME classification bands in tool choice and factor formality and level.

Third, it adds tentative indications of factors becoming increasingly relevant in later-stage adoption. For instance, facilitating conditions were found not to be relevant in early-stage adoption (Kim, Blazquez, and Oh 2024) whereas in the later-stage adoption conditions of this study they are increasingly mentioned. Additionally, regulatory concerns initially were contested (Hughes et al. 2026; L. Zhang et al. 2025), however, they appear to be more pertinent with the passing of time than in earlier studies.

5.4 Implications for Practice

This study proposes three implications for small to medium enterprises adopting generative AI as well as a few for tool vendors providing the tools.

For SMEs

First, provisioning is a strong lever to pull to positively influence adoption, where paying for one tool is already beneficial as it lowers the barrier to adopt the technology. Direct provisioning of one or multiple subscriptions is the most effective choice, followed by allowing for personal accounts to be used. Shared and free accounts generally should be the last resort as both have inherent limitations that, at times, lead employees to disregard policy and use personal accounts instead.

Second, a permissive-default environment combined with frequent alignment moments on unwritten norms beats restrictive policy which may be disregarded and unrealistic to enforce. That is not to say that policy is not effective, however it should cover both restrictions to constrain and possibilities to enable adoption.

Third, budgets are finite and choices need to be made on who gets what subscription. This choice should not hinge on the role of the employee but on the expected return in terms of efficiency gains, output quality or being able to do something that previously was not possible due to skill limitations. Direct managers have the clearest perspective to make such choices.

For tool vendors

This study implies that task-tool fit directly shapes which tools are adopted. The importance of training the models on concrete enterprise tasks is crucial, that is if the vendor targets enterprise customers. Specialised spin-offs of general-purpose tools, such as Claude Cowork, are especially effective where the specialised tool already is associated with a habitually-used and trusted general-purpose tool's brand name. For specialised tool vendors, it is advised to aim for a niche capability which is not well supported by general-purpose tools and cannot easily be replicated or is not worth replicating by established vendors.

5.5 Limitations

Several limitations bound the interpretation of these findings. By design, the study favoured breadth over depth. It surfaced tendencies, yet not statistically valid relationships between organisational and individual factors. It enumerated the factors at play. However, their effects on one another and on behavioural intention to adopt generative AI are beyond its scope. A number of factors rested on only a few unique mentions and may not generalise beyond the firms that raised them, where some likely reflect the organisation's particular context. The study, therefore, offers a basis for deeper, more targeted research rather than claiming depth itself.

The sample size was formed by convenience sampling, which limits how far findings generalise. Nineteen interviews may not fully saturate every theme, and the size bands are especially thin consisting of 4 micro and 5 small enterprises as compared to 10 medium ones. So, factors that specifically relate to the smaller firms may be under-represented. While the sample spans many industries, it cannot capture the full diversity of the SME population. These constraints affect which factors surface and how often, but the relative appearance of factors and the way they are clustered offer a credible picture of how adoption occurs.

Nevertheless, this study provides an empirically grounded account of generative AI adoption in small to medium enterprises that bridges the gap between the individual and organisational levels usually examined in isolation.

5.6 Future Research

Future research could concentrate on quantitatively validating the tendencies this study surfaced. The identified factors could serve as a basis of a survey instrument, measuring their influence on adoption. This study delivers a UTAUT-based survey instrument to investigate individual generative AI adoption as shaped by organisational dynamics in Appendix C.

Longitudinal studies could track how adoption of generative AI evolves within a single firm, testing the temporal patterns which this study could only infer. In particular, whether informal practice precedes formal policy and how governance practices adapt to facilitate adoption of emerging generative AI capabilities. This would clarify the regulatory-practice transition this study proposed but could not observe over time.

This study could be replicated across geographical, cultural and demographical contexts, with a larger sample size. This study suggests that social influence operates laterally and not by rank, but inferred that it may not hold in a higher power-distance environment. Cross-cultural replication could test that boundary directly. A larger, more balanced spread of firm sizes would allow micro, small and medium companies to be examined more deeply. This study indicated that influencing factors differ by size but on bands too thin to generalise. Partitioning by industry, rather than size alone, could further reveal sectoral differences in adoption practice.

Chapter 6

Conclusions

This research set out to examine how individual and organisational factors affect generative AI adoption in small to medium enterprises, guided by the central question: *“How do SMEs overcome barriers and create enablers to increase generative AI adoption?”*. It finds that the adoption of generative AI is bottom-up, individual-led, with the organisation enabling and reacting rather than initiating. Adopters favour general-purpose tools, though specific reasons lead to specialised tools to be adopted in particular cases. Governance is permissive by default, with formal policy the exception and unwritten norms the rule. Notably, the formality of factors tends to increase with firm size, where such factors likewise become progressively organisational.

Where previous literature treated individual and organisational adoption factors separately, this research shows they are, in fact, two sides of the same coin in the SME context. It demonstrates how organisational decisions—relating to tool provisioning, risk management and organisational culture, among others—act through individuals, shaping organisational adoption by influencing individuals’ choices. This equips SMEs with a clearer understanding of the generative AI adoption dynamic, allowing them to navigate uncertain waters with greater confidence in the future.

Beyond any single lever, this study reframes how adoption in SMEs should be understood. It is driven by individuals adopting the technology, and predominantly not anything initiated by the organisation. This leaves the firms to enable and govern after adoption has taken place. As these tools grow more capable and increasingly autonomous, the bottom-up dynamic observed is likely to deepen, and the tension between permissive adoption and risk-driven governance that firms are only beginning to navigate will progressively become increasingly complicated and challenging.

This raises the question that contemporary adoption literature is not yet capable of answering: *when a technology is adopted from the bottom-up, how should its governance be designed to strike a balance between enabling and restricting adoption?* Readers should now see that, in the SME setting, the individual and organisation cannot be studied apart. Each shapes the other, and adoption is the product of their interaction.

Bibliography

- Ajzen, Icek (1991). “The Theory of Planned Behavior”. In: *Organizational Behavior and Human Decision Processes* 50.2, pp. 179–211. DOI: [10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T).
- Baker, Jeff (2012). “The Technology–Organization–Environment Framework”. In: *Information Systems Theory: Explaining and Predicting Our Digital Society, Vol. 1*. Ed. by Yogesh K. Dwivedi, Michael R. Wade, and Scott L. Schneberger. New York: Springer, pp. 231–245. DOI: [10.1007/978-1-4419-6108-2_12](https://doi.org/10.1007/978-1-4419-6108-2_12).
- Bengio, Yoshua et al. (2003). “A Neural Probabilistic Language Model”. In: *Journal of Machine Learning Research* 3, pp. 1137–1155.
- Bick, Alexander, Adam Blandin, and David J. Deming (2026). “The Rapid Adoption of Generative AI”. In: *Management Science*. DOI: [10.1287/mnsc.2025.02523](https://doi.org/10.1287/mnsc.2025.02523). URL: <https://doi.org/10.1287/mnsc.2025.02523> (visited on 06/28/2026).
- Biloš, Antun and Bruno Budimir (2024). “Understanding the Adoption Dynamics of ChatGPT among Generation Z: Insights from a Modified UTAUT2 Model”. In: *Journal of Theoretical and Applied Electronic Commerce Research* 19.2, pp. 863–879. DOI: [10.3390/jtaer19020045](https://doi.org/10.3390/jtaer19020045).
- Boonmee, Chawis, Jettapat Mangkalakeeree, and Yongkuk Jeong (2025). “Towards sustainable digital transformation: AI adoption barriers and enablers among SMEs in Northern Thailand”. In: *Sustainable Futures* 10, p. 101169. ISSN: 2666-1888. DOI: [10.1016/j.sftr.2025.101169](https://doi.org/10.1016/j.sftr.2025.101169).
- Braun, Virginia and Victoria Clarke (2006). “Using Thematic Analysis in Psychology”. In: *Qualitative Research in Psychology* 3.2, pp. 77–101. DOI: [10.1191/1478088706qp063oa](https://doi.org/10.1191/1478088706qp063oa).
- Brown, Tom B. et al. (2020). “Language Models are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*.
- Camilleri, Mark Anthony (2024). “Factors Affecting Performance Expectancy and Intentions to Use ChatGPT: Using SmartPLS to Advance an Information Technology Acceptance Framework”. In: *Technological Forecasting and Social Change* 201, p. 123247. DOI: [10.1016/j.techfore.2024.123247](https://doi.org/10.1016/j.techfore.2024.123247).
- Davis, Fred D. (1989). “Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology”. In: *MIS Quarterly* 13.3, pp. 319–340. DOI: [10.2307/249008](https://doi.org/10.2307/249008).
- DiMaggio, Paul J. and Walter W. Powell (1983). “The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields”. In: *American Sociological Review* 48.2, pp. 147–160. DOI: [10.2307/2095101](https://doi.org/10.2307/2095101).
- Eisfeldt, Andrea L. and Gregor Schubert (Oct. 2024). “Generative AI and Finance”. In: SSRN working paper. DOI: [10.2139/ssrn.4988553](https://doi.org/10.2139/ssrn.4988553). URL: <https://ssrn.com/abstract=4988553>.
- Elman, Jeffrey L. (1990). “Finding structure in time”. In: *Cognitive Science* 14.2, pp. 179–211.
- Fang, Cong et al. (2026). “A Comparative Analysis of Generative AI Adoption among Design Professionals in China and the United Kingdom: A UTAUT Perspective”. In: *Humanities and Social Sciences Communications* 13.1, p. 411. DOI: [10.1057/s41599-026-06796-x](https://doi.org/10.1057/s41599-026-06796-x).

- Felder, Ludwig et al. (Jan. 2026). “Adoption of Generative Artificial Intelligence in the German Software Engineering Industry: An Empirical Study”. In: *arXiv e-prints*, arXiv:2601.16700, arXiv:2601.16700. DOI: [10.48550/arXiv.2601.16700](https://doi.org/10.48550/arXiv.2601.16700). arXiv: [2601.16700](https://arxiv.org/abs/2601.16700) [cs.SE].
- Fishbein, Martin and Icek Ajzen (1975). *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT Press.
- Goodfellow, Ian, Jean Pouget-Abadie, et al. (2014). “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems*.
- Halevy, Alon, Peter Norvig, and Fernando Pereira (Mar. 2009). “The Unreasonable Effectiveness of Data”. In: *IEEE Intelligent Systems* 24.2, pp. 8–12. ISSN: 1541-1672. DOI: [10.1109/MIS.2009.36](https://doi.org/10.1109/MIS.2009.36). URL: <https://doi.org/10.1109/MIS.2009.36>.
- He, Chen et al. (2026). “Navigating the Generative Frontier: Adoption, Challenges, and Strategic Realignment of GAI in the Chinese Publishing Industry”. In: *Information Research* 31.iConf, pp. 606–626. DOI: [10.47989/ir31iConf64142](https://doi.org/10.47989/ir31iConf64142).
- Ho, Jonathan, Ajay Jain, and Pieter Abbeel (2020). “Denoising Diffusion Probabilistic Models”. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
- Hochreiter, Sepp and Jürgen Schmidhuber (1997). “Long Short-Term Memory”. In: *Neural Computation* 9.8, pp. 1735–1780.
- Hu, Krystal (Feb. 2023). *ChatGPT Sets Record for Fastest-Growing User Base — Analyst Note*. Reuters. URL: <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/> (visited on 06/28/2026).
- Hughes, Laurie et al. (2026). “Beyond the hype: Organisational adoption of Generative AI through the lens of the TOE framework—A mixed methods perspective”. In: *International Journal of Information Management* 86, p. 102982. DOI: [10.1016/j.ijinfomgt.2025.102982](https://doi.org/10.1016/j.ijinfomgt.2025.102982).
- Kaplan, Jared et al. (2020). “Scaling Laws for Neural Language Models”. In: *arXiv preprint arXiv:2001.08361*.
- Kim, Youngsoo, Victor Blazquez, and Taeyeon Oh (2024). “Determinants of Generative AI System Adoption and Usage Behavior in Korean Companies: Applying the UTAUT Model”. In: *Behavioral Sciences* 14.11, p. 1035. DOI: [10.3390/bs14111035](https://doi.org/10.3390/bs14111035).
- Kipf, Thomas N. and Max Welling (2017). “Semi-Supervised Classification with Graph Convolutional Networks”. In: *International Conference on Learning Representations*.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey Hinton (2012). “ImageNet Classification with Deep Convolutional Neural Networks”. In: *NeurIPS*.
- Kumar, Aman et al. (2025). “Generative Artificial Intelligence (GenAI) Revolution: A Deep Dive into GenAI Adoption”. In: *Journal of Business Research* 189, p. 115160. ISSN: 0148-2963. DOI: [10.1016/j.jbusres.2024.115160](https://doi.org/10.1016/j.jbusres.2024.115160).
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015). “Deep Learning”. In: *Nature* 521, pp. 436–444.
- LeCun, Yann, Léon Bottou, et al. (1998). “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11, pp. 2278–2324.
- Lim, Rachel Esther et al. (2025). “Excited or Anxious about Adopting Generative AI? Investigating the Drivers of Generative AI Adoption among Advertising Creatives”. In: *Journal of Marketing Communications*. Advance online publication. DOI: [10.1080/13527266.2025.2497259](https://doi.org/10.1080/13527266.2025.2497259).
- Looi, Mark (2026). “Developers in the Age of AI: Adoption, Policy, and Diffusion of AI Software Engineering Tools”. In: *CoRR* abs/2601.21305. DOI: [10.48550/arXiv.2601.21305](https://doi.org/10.48550/arXiv.2601.21305).
- Metzger, Michael, Seán O’Reilly, and Ciarán Mac an Bhaird (2025). “Generative Artificial Intelligence Augmenting SME Financial Management”. In: *Technovation* 147, p. 103313. ISSN: 0166-4972. DOI: [10.1016/j.technovation.2025.103313](https://doi.org/10.1016/j.technovation.2025.103313).
- Mitchell, Tom M. (1997). *Machine Learning*. McGraw-Hill.

- Parasuraman, A. (2000). "Technology Readiness Index (TRI): A Multiple-Item Scale to Measure Readiness to Embrace New Technologies". In: *Journal of Service Research* 2.4, pp. 307–320. DOI: [10.1177/109467050024001](https://doi.org/10.1177/109467050024001).
- Parasuraman, A. and Charles L. Colby (2015). "An Updated and Streamlined Technology Readiness Index: TRI 2.0". In: *Journal of Service Research* 18.1, pp. 59–74. DOI: [10.1177/1094670514539730](https://doi.org/10.1177/1094670514539730).
- Patel, Kinjal et al. (2026). "A systematic review of generative AI: importance of industry and startup-centered perspectives, agentic AI, ethical considerations & challenges, and future directions". In: *Artificial Intelligence Review* 59.1, p. 7. DOI: [10.1007/s10462-025-11435-z](https://doi.org/10.1007/s10462-025-11435-z). URL: <https://doi.org/10.1007/s10462-025-11435-z>.
- Radford, Alec, Karthik Narasimhan, et al. (2018). "Improving Language Understanding by Generative Pre-Training". In: *OpenAI Technical Report*.
- Radford, Alec, Jeffrey Wu, et al. (2019). "Language Models are Unsupervised Multitask Learners". In: *OpenAI Technical Report*.
- Rijksdienst voor Ondernemend Nederland (June 10, 2025). *Mkb-verklaring*. Accessed 24 February 2026. Ministerie van Economische Zaken. URL: <https://www.rvo.nl/onderwerpen/subsidiespelregels/emkb-verklaring> (visited on 02/24/2026).
- Rogers, Everett M. (2003). *Diffusion of Innovations*. 5th ed. New York: Free Press.
- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams (1986). "Learning representations by back-propagating errors". In: *Nature* 323, pp. 533–536.
- Russell, Stuart and Peter Norvig (2010). *Artificial Intelligence: A Modern Approach*. Pearson.
- Salih, Sayeed et al. (2026). "From ideation to execution: Unleashing the power of generative AI in modern digital marketing and customer engagement — A systematic literature review and case study". In: *Array* 29, p. 100630. ISSN: 2590-0056. DOI: [10.1016/j.array.2025.100630](https://doi.org/10.1016/j.array.2025.100630).
- Sharma, A., P. Kakkar, S. Hooda, et al. (2026). "How Does Gen-AI Reshape Supply Chain, Logistics, and Manufacturing? A Systematic Bibliometric Review". In: *Operations Research Forum* 7.1, p. 5. DOI: [10.1007/s43069-025-00582-2](https://doi.org/10.1007/s43069-025-00582-2).
- Smutny, Zdenek and Frantisek Sudzina (2025). "What Affects Work Performance When Using AI Chatbots? Investigating Mediations and Factors Affecting Performance Expectancy and Intentions to Use ChatGPT". In: *International Journal of Human-Computer Interaction*. Advance online publication. DOI: [10.1080/10447318.2025.2573037](https://doi.org/10.1080/10447318.2025.2573037).
- Stanford Institute for Human-Centered Artificial Intelligence (2026). *The 2026 AI Index Report*. Stanford University. URL: <https://hai.stanford.edu/ai-index/2026-ai-index-report> (visited on 06/28/2026).
- Taiwo, Ridwan et al. (2025). "Generative artificial intelligence in construction: A Delphi approach, framework, and case study". In: *Alexandria Engineering Journal* 116, pp. 672–698. ISSN: 1110-0168. DOI: [10.1016/j.aej.2024.12.079](https://doi.org/10.1016/j.aej.2024.12.079).
- Tewari, Veena, Amal Sulaiman Al Alawi, and Swapnil Morande (2026). "Generative Artificial Intelligence Adoption in Emerging Economies: A Technology–Organization–Environment Framework Analysis of Large Language Model Integration in Small and Medium Enterprises". In: *International Journal of Research and Innovation in Social Science* 10.19. DOI: [10.47772/IJRIS.2026.10190045](https://doi.org/10.47772/IJRIS.2026.10190045).
- Thong, James Y. L. (1999). "An Integrated Model of Information Systems Adoption in Small Businesses". In: *Journal of Management Information Systems* 15.4, pp. 187–214. DOI: [10.1080/07421222.1999.11518227](https://doi.org/10.1080/07421222.1999.11518227).
- Tornatzky, Louis G. and Mitchell Fleischer (1990). *The Processes of Technological Innovation*. Lexington, MA: Lexington Books.
- Vaswani, Ashish et al. (2017). "Attention Is All You Need". In: *Advances in Neural Information Processing Systems (NeurIPS)*.

- Venkatesh, Viswanath and Hillol Bala (2008). "Technology Acceptance Model 3 and a Research Agenda on Interventions". In: *Decision Sciences* 39.2, pp. 273–315. DOI: [10.1111/j.1540-5915.2008.00192.x](https://doi.org/10.1111/j.1540-5915.2008.00192.x).
- Venkatesh, Viswanath and Fred D. Davis (2000). "A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies". In: *Management Science* 46.2, pp. 186–204. DOI: [10.1287/mnsc.46.2.186.11926](https://doi.org/10.1287/mnsc.46.2.186.11926).
- Venkatesh, Viswanath, Michael G. Morris, et al. (2003). "User Acceptance of Information Technology: Toward a Unified View". In: *MIS Quarterly* 27.3, pp. 425–478. DOI: [10.2307/30036540](https://doi.org/10.2307/30036540).
- Woo, Sang Eun, Ernest H. O’Boyle, and Paul E. Spector (2017). "Best Practices in Developing, Conducting, and Evaluating Inductive Research". In: *Human Resource Management Review* 27.2, pp. 255–264. ISSN: 1053-4822. DOI: [10.1016/j.hrmr.2016.08.004](https://doi.org/10.1016/j.hrmr.2016.08.004).
- Xia, Yan and Yue Chen (2025). "Driving Factors of Generative AI Adoption in New Product Development Teams from a UTAUT Perspective". In: *International Journal of Human–Computer Interaction* 41.10, pp. 6067–6088. DOI: [10.1080/10447318.2024.2375686](https://doi.org/10.1080/10447318.2024.2375686).
- Xin, Baogui, Chunyu Zhang, and Wei Peng (2025). "When Technical Literacy Backfires: Rethinking Employee Acceptance of Generative AI via Extended UTAUT". In: *International Journal of Human–Computer Interaction*. Advance online publication. DOI: [10.1080/10447318.2025.2565399](https://doi.org/10.1080/10447318.2025.2565399).
- Xu, Zaijuan, T. Ramayah, and Yubo Shi (2025). "Determinants of Generative Artificial Intelligence Adoption in Insurance Companies". In: *International Journal of Business and Society* 26.1, pp. 272–291. DOI: [10.33736/ijbs.9562.2025](https://doi.org/10.33736/ijbs.9562.2025).
- Yesuf, Y. and Z. Fields (2025). "Artificial Intelligence Adoption as a Driver of Innovation and Competitiveness in SMEs: A Bibliometric and Systematic Review". In: *F1000Research* 14, p. 1187. DOI: [10.12688/f1000research.171494.1](https://doi.org/10.12688/f1000research.171494.1).
- Yesuf, Yohannes Mekonnen and Ziska Fields (Oct. 2025). "Artificial Intelligence Adoption as a Driver of Innovation and Competitiveness in SMEs: A Bibliometric and Systematic Review". In: *F1000Research* 14, p. 1187. DOI: [10.12688/f1000research.171494.1](https://doi.org/10.12688/f1000research.171494.1).
- Zhang, L. et al. (2025). "Unraveling Generative AI Adoption in Enterprise Digital Platforms: The Effect of Institutional Pressures and the Moderating Role of Internal and External Environments". In: *IEEE Transactions on Engineering Management* 72, pp. 335–348. DOI: [10.1109/TEM.2024.3513773](https://doi.org/10.1109/TEM.2024.3513773).

Appendix A

Consent Form

Interview Consent Form

This interview is conducted by Jakub Szymon Korus, a student in the Informatica & Economie programme at Leiden University, as part of a bachelor's thesis project supervised by T.D. Offerman MSc. It is part of a bachelor's thesis on how small and medium-sized enterprises adopt generative AI, and what organisational factors influence this.

Audio recording

This interview will be audio-recorded. The recording is used solely for transcription purposes and will not be shared with anyone outside the research team.

What is recorded

Your verbal statements during the interview, and video footage if you choose to enable your camera.

Anonymity

Everything you say will be anonymised during transcription. No statements will be traceable back to you.

Data storage

Data is stored securely on the interviewer's personal device and Leiden University's OneDrive, and deleted upon completion of the thesis.

Voluntary participation

Participation is voluntary. You may withdraw consent at any point until the interviewer declares the interview concluded. After that point, deletion can no longer be requested, as the data will be used for academic research purposes.

Consent

By confirming verbally or in writing, you confirm that you have read this form and agree to participate.

Questions? Contact me at j.s.korus@umail.leidenuniv.nl

Appendix B

Interview Protocol

Table B.1. Semi-structured interview guide

Introduction	
<i>Procedural: recording and anonymity note, consent form, personal introduction, research aim and question.</i>	
Generative AI in their job	
1	How do you use Generative AI in your job?
2	Was the choice to use it your own, or was it encouraged or mandated by the company?
3	Which specific tools?
4	How often do you use it? Would you say you rely on it?
5	Which tools do others at your company use, and for what?
Company's Perspective	
6	What is your company's stance on Generative AI use?
7	How do you feel about that stance as an employee?
Enablers and Barriers	
8	What has made it easier for you to use Generative AI at work?
9	What has made it harder?
10	What formal things has your company provided or put in place around Generative AI?
11	Have you ever been technically or practically blocked from using a tool? If so, in what context?
Roles & Social Influence	
12	What informal things have shaped how you use Generative AI at work?
13	Which roles at your company have made it easier for you to use Generative AI?
14	Which roles have made it harder?
15	Do your direct colleagues influence how you use it?
Closing	
16	Is there anything we haven't covered that you think is relevant?

Appendix C

Proposed Survey Instrument

Table C.1. Proposed UTAUT-based survey instrument

Background			
<i>How many employees does your company have?</i>			
Micro (1–10) / Small (11–50) / Medium (51–250)			
<i>How long has generative AI been in regular use at your organisation?</i>			
Under 6 months / 6–18 months / Over 18 months			
Construct	Code	Item	Source
Performance Expectancy	PE1	Generative AI helps complete tasks more easily.	(Kim, Blazquez, and Oh 2024)
	PE2	Generative AI contributes to improving work performance and efficiency.	(Kim, Blazquez, and Oh 2024)
	PE3	Generative AI helps solve complex problems.	(Kim, Blazquez, and Oh 2024)
Perceived Risk	PR1	I have privacy concerns when using Generative AI.	(Xia and Chen 2025)
	PR2	I am concerned that using Generative AI may bring reputational risk to my company.	This study
	PR3	I am concerned that using Generative AI may make me too reliant on it.	(Xia and Chen 2025)
Effort Expectancy	EE1	Learning how to use Generative AI is easy for me.	(Xia and Chen 2025)
	EE2	My interaction with Generative AI is clear and understandable.	(Xia and Chen 2025)
	EE3	Using Generative AI is a convenient and efficient process that does not require a significant amount of time.	(Xia and Chen 2025)
Social Influence	SI1	Seeing colleagues use generative AI influences how I use it.	This study
	SI2	Colleagues who are skilled with generative AI influence how I use it.	This study
	SI3	My manager’s stance on generative AI influences how I use it.	This study

Table C.1 (continued)

Construct	Code	Item	Source
Facilitating Conditions	FC1	Top management actively supports the adoption and use of generative AI.	(Kim, Blazquez, and Oh 2024)
	FC2	My company provides the generative AI subscriptions I need for my work.	This study
	FC3	There is sufficient support from AI experts within the organisation for the use of generative AI.	(Kim, Blazquez, and Oh 2024)
	FC4	My company provides sufficient training for the use of generative AI.	(Kim, Blazquez, and Oh 2024)
Behavioural Intention	BI1	I intend to continue using generative AI at work in the future.	(Fang et al. 2026)
	BI2	I will always try to use generative AI in my daily work.	(Fang et al. 2026)
Usage Behaviour	UB1	I personally use generative AI like ChatGPT.	(Kim, Blazquez, and Oh 2024)
	UB2	I use generative AI to try new ways of working.	(Kim, Blazquez, and Oh 2024)