



AI Maturity and Workforce Displacement: A Qualitative Comparative Study of Governance Across Maturity Stages

Yannick Schulz

Dissertation written under the supervision of Professor René Bohnsack

Dissertation submitted in partial fulfilment of requirements for the MSc in Management with
a specialization in Strategy, Entrepreneurship & Impact, at the Universidade Católica
Portuguesa, 28.05.2026.

Author: Yannick Schulz

Title: AI Maturity and Workforce Displacement: A Qualitative Comparative Study of Governance Across Maturity Stages

Submission Date: 28 May 2026

Key Words: Artificial intelligence; AI maturity; workforce displacement; AI governance; task reallocation

Abstract

Artificial intelligence (AI) is increasingly embedded in organisational activities, reshaping how tasks are performed, coordinated, and governed. While research on AI-related labour-market effects has focused mainly on occupational or sectoral exposure, AI maturity research has examined capability development, and AI governance research has rarely connected governance practices to specific displacement pressures. Less is known about how organisations at different AI maturity stages experience and govern workforce displacement. This study addresses this gap by examining AI-enabled workforce displacement at the firm level.

The study draws on thirteen semi-structured expert interviews with firm-side and cross-organisational informants. Transcripts were analysed through a Gioia-inspired abductive coding procedure.

The findings show that AI implementation differed not only in tool availability, but in how AI was coordinated and embedded. AI capability remained unevenly distributed across functions, sites, and employee groups. Workforce displacement appeared most visibly as task reallocation within existing roles; role elimination and headcount reduction remained limited, indirect, or anticipated. The conversion of technical capability into realised workforce effects was shaped by organisational, legal, technical, and institutional filters. Governance developed unevenly: mitigation, the organisation of adaptation, became more formalised, while legitimisation, the explanation of workforce consequences, remained indirect.

The study contributes a firm-level account of AI-enabled workforce displacement as a stage-conditioned conversion process rather than a direct result of occupational exposure. It further shows that AI maturity should be read as a dominant organisational configuration, and that responsible AI governance should be assessed by whether organisations communicate workforce consequences in explainable, contestable, and justifiable ways.

Autor: Yannick Schulz

Título: Maturidade em IA e Deslocamento da Força de Trabalho: Um Estudo Qualitativo Comparativo sobre a Governança em Diferentes Estágios de Maturidade

Data de envio: 28 de maio de 2026

Palavras-chave: Inteligência artificial; maturidade em IA; deslocamento da força de trabalho; governança da IA; realocação de tarefas

Resumo

A inteligência artificial (IA) está crescentemente integrada nas atividades organizacionais, reconfigurando tarefas, coordenação e governança. Embora os estudos sobre efeitos laborais da IA se concentrem na exposição ocupacional ou setorial, a investigação sobre maturidade analisa capacidades, e a investigação sobre governança raramente liga práticas a pressões de deslocamento. Sabe-se menos sobre como organizações em diferentes estágios de maturidade experienciam e governam o deslocamento laboral. Este estudo responde a esta lacuna ao examinar este fenómeno ao nível da firma.

O estudo baseia-se em treze entrevistas semiestruturadas com informantes internos e transorganizacionais, analisadas por codificação abductiva inspirada em Gioia.

Os resultados mostram que a implementação da IA diferiu na disponibilidade de ferramentas, coordenação e integração. A capacidade de IA permaneceu desigual entre funções, locais e grupos de trabalhadores. O deslocamento laboral surgiu sobretudo como realocação de tarefas; a eliminação de funções e a redução de efetivos permaneceram limitadas, indiretas ou antecipadas. A conversão de capacidade técnica em efeitos laborais realizados foi moldada por filtros organizacionais, legais, técnicos e institucionais. A governança desenvolveu-se de forma desigual: a mitigação tornou-se mais formalizada, enquanto a legitimação permaneceu indireta.

O estudo contribui com uma explicação ao nível da firma do deslocamento laboral possibilitado pela IA como processo de conversão condicionado pela maturidade, e não como resultado direto da exposição ocupacional. Mostra ainda que a maturidade em IA deve ser lida como configuração organizacional dominante, e que a governança responsável da IA deve ser avaliada pela comunicação de consequências laborais de forma explicável, contestável e justificável.

Table of Contents

1	Introduction	1
2	Literature Review	3
2.1.	Artificial Intelligence as a Firm-Level Phenomenon	3
2.1.1.	Defining AI for Firm-Level Analysis	3
2.1.2.	AI Deployment as a Firm-Level Process	4
2.2.	AI Maturity as an Analytical Framework	5
2.3.	Workforce Displacement as an Organisational Outcome	7
2.3.1.	Defining Workforce Displacement at the Firm Level	7
2.3.2.	From Exposure Estimates to Observed Workforce Effects	9
2.3.3.	Variation Across AI Maturity Stages	10
2.4.	Ethics and Governance of AI-Enabled Workforce Displacement	10
2.4.1.	Ethical Standards for AI-Enabled Workforce Displacement	10
2.4.2.	Governance as Mitigation	12
2.4.3.	Governance as Legitimation	13
2.5.	Integrated Conceptual Framing	14
3	Methodology	16
3.1.	Research Design	16
3.2.	Data Collection	16
3.3.	Data Analysis	19
4	Findings	20
4.1.	AI Maturity as Formal Coordination with Internal Differences	22
4.2.	Workforce Displacement Concentrated at Task Level	24
4.3.	Governance Visible Mainly as Mitigation	27
5	Discussion	29
5.1.	AI Maturity as a Dominant Organisational Configuration	30
5.2.	Workforce Displacement as a Stage-Conditioned Conversion Process	31
5.3.	Governance as Mitigation-Legitimation Asymmetry	33
5.4.	Integrated Framework	35
6	Conclusion and Outlook	36
6.1.	Concluding Answers to the Research Questions	36
6.2.	Theoretical Contributions and Practical Implications	38
6.3.	Limitations, Future Research, and Closing Reflection	39
	References	42
	Appendix	46

List of Figures

Figure 1: Maturity dimension terminology across three frameworks	6
Figure 2: Analytical expectations across AI maturity stages	15
Figure 3: AI maturity classification of firm-side informants	18
Figure 4: Focus area and maturity range of cross-organisational informants	19
Figure 5: Data structure of the empirical analysis.....	21
Figure 6: Stage-sensitive and stage-qualifying empirical patterns.....	29
Figure 7: Empirical refinements to the framework logic	36

List of Abbreviations

Abbreviation	Meaning
AI	Artificial Intelligence
CV	Curriculum Vitae
EU	European Union
HR	Human Resources
LLM	Large Language Model
RAG	Retrieval-Augmented Generation
RQ	Research Question
U.S.	United States

Acknowledgements

First, I would like to thank my supervisor, Mickie de Wet, for her continuous guidance, thoughtful feedback, and support throughout the process of writing this thesis. Her comments helped me to develop my ideas more clearly and to improve the overall quality of this work. I would also like to thank the interview partners who participated in this study. Without their time, openness, and valuable insights, this thesis would not have been possible.

A special thank you goes to my friends with whom I spent a month in Lisbon while we were all working on our master's theses. Sharing this time made the writing process not only much more enjoyable but also enriched this thesis through many conversations and different perspectives on the topic.

I also want to thank my girlfriend, who not only read my thesis and provided valuable feedback but also supported me throughout the entire process.

Finally, I would like to thank my parents. Their continuous support, trust, and encouragement made it possible for me to pursue this master's degree in the first place. I am deeply grateful for everything they have done for me, not only during this thesis, but throughout my entire academic path.

1 Introduction

Artificial intelligence (AI) is becoming embedded in organisational operations across a widening range of activities. Recent estimates suggest that around 80% of the U.S. workforce could see at least 10% of their tasks affected by large language models (LLMs) alone (Eloundou et al., 2024). Yet interpreting AI primarily as a force of job elimination misrepresents how it changes work in organisational settings. Rather than replacing entire positions, AI more often reconfigures the task composition of roles and occupations (Acemoglu & Restrepo, 2019; Agrawal et al., 2019; Bankins et al., 2024). Whether a given application displaces labour, augments human work, or generates new tasks depends on the tasks involved, the organisational context, and the conditions of deployment (Clifton et al., 2020; Felten et al., 2021).

Three insights from prior research explain why these workforce effects are difficult to predict. First, AI changes work at the level of tasks and decisions, because its applications reduce the cost of prediction within organisational processes. As prediction becomes cheaper, other components of decision-making, especially judgement, interpretation, and action, become relatively more important, which means that some task components become automatable while others remain human. For instance, an AI system may automate contract review while lawyers retain responsibility for interpreting findings and advising clients (Agrawal et al., 2019). Second, automation does not only displace labour. It can also trigger reinstatement through new tasks in which labour retains a comparative advantage (Acemoglu & Restrepo, 2019). Workforce effects therefore emerge through the balance between displacement and reinstatement rather than through substitution alone, and that balance varies across industries, skill levels, and time horizons (Raisch & Krakowski, 2021). Third, at the firm level, automation and augmentation are not stable alternatives but often evolve in a cyclical relationship over time (Raisch & Krakowski, 2021). Tasks initially performed in collaboration with AI may later become automated, while new adjacent tasks emerge for human workers.

Despite this, the relevant literatures do not yet provide a firm-level explanation of how these pressures vary across organisations. Research on the labour market effects of AI has examined task exposure, displacement, and reinstatement, but it predominantly operates at the occupational or sectoral level and therefore says little about how displacement pressures differ across firms (Felten et al., 2021). Research on AI maturity explains how organisations develop AI capabilities across dimensions such as data, technology, skills, integration, and strategic

alignment, but it focuses primarily on readiness and capability development rather than workforce consequences (Sadiq et al., 2021). Research on AI governance and ethics addresses responsible deployment, yet rarely examines how governance mechanisms are used to manage, mitigate, or legitimate workforce displacement as AI becomes more deeply embedded in the firm (Papagiannidis et al., 2025). This leaves a gap in explaining how workforce displacement varies across organisations and how firms govern these pressures at different stages of AI maturity.

This study addresses this gap by developing an exploratory conceptual framework that links AI maturity stages to distinct patterns of workforce displacement and to the governance mechanisms firms use in response. The study focuses on the firm level because displacement is shaped by organisational choices about where AI is deployed, how tasks are redesigned, and how workforce pressures are managed, variation that occupational or sectoral analyses cannot capture (Clifton et al., 2020; Raisch & Krakowski, 2021). Empirically, the study adopts a qualitative comparative expert-interview design that draws on informed accounts of large enterprises at different stages of AI maturity. The resulting framework is exploratory and intended for further empirical refinement. The study is guided by the following research question:

RQ: How do organisations at different AI maturity stages experience and govern workforce displacement?

This overarching question is addressed through two subordinate questions:

(1) What workforce displacement patterns are associated with different AI maturity stages?

(2) Which governance mechanisms do firms use to mitigate and legitimate displacement across maturity stages?

By connecting research on AI maturity, workforce displacement and governance, the study makes both a theoretical and a practical contribution. Theoretically, it introduces AI maturity as a conditioning factor for workforce displacement at the firm level and examines how governance responses differ across stages of AI integration. Practically, the framework can help organisations anticipate the displacement pressures likely to emerge at their current stage of AI adoption and identify relevant governance responses.

The study proceeds as follows. Chapter 2 reviews the literature on AI as a firm-level phenomenon, AI maturity, workforce displacement, and the ethics and governance of AI-enabled workforce change, before developing an integrated conceptual framing. Chapter 3 presents the qualitative comparative methodology, including the interview design, sampling strategy, and analytical approach. Chapter 4 reports the empirical findings across the maturity-stage contexts identified in the data. Chapter 5 discusses the results and develops the integrated framework. Chapter 6 concludes by answering the research questions, outlining theoretical and practical contributions, acknowledging limitations, and identifying directions for future research.

2 Literature Review

In this chapter, the conceptual basis for the study is developed by reviewing AI as a firm-level phenomenon, AI maturity models, workforce displacement, and the governance of AI-enabled workforce change.

2.1. Artificial Intelligence as a Firm-Level Phenomenon

2.1.1. Defining AI for Firm-Level Analysis

A firm-level analysis of AI-related workforce consequences requires a clear conceptualisation of what is meant by AI. The term AI is commonly used as an umbrella label for a broad set of computational technologies, including machine learning, deep learning, natural language processing, computer vision, and generative models (Sadiq et al., 2021; Bankins et al., 2024). At the same time, the terms AI, machine learning, and automation are often used interchangeably in academic and practitioner discourse, while managers may classify routine automation as AI even when the underlying technology involves limited machine intelligence (Clifton et al., 2020).

Against this background, several definitions in the literature adopt a functional rather than a purely technical perspective. Bankins et al. (2024) define AI as a collection of interrelated technologies used to solve problems and perform tasks that, when carried out by humans, require thinking. Similarly, Krakowski et al. (2023) describe AI as enabling machines to perform cognitive functions previously associated with human minds. These definitions focus less on technical architecture and more on what AI systems do within organisational settings.

The present study adopts the definition by Bankins et al. (2024), who describe AI as a collection of interrelated technologies used to solve problems and perform tasks that, when carried out by humans, require thinking. Its functional and task-based framing is better suited to a firm-level analysis of workforce consequences than a purely technical perspective, because workforce effects depend on what AI does in work processes rather than on the technologies involved. This definition can therefore include technologies such as machine learning, deep learning, natural language processing, computer vision, and generative models, but is not limited to them.

2.1.2. AI Deployment as a Firm-Level Process

With this functional definition established, the next step is to explain why AI must be analysed at firm level. Mikalef and Gupta (2021) conceptualise AI capability as a firm's ability to select, orchestrate, and leverage AI-specific resources across three categories: tangible resources, including data, technology, and budget, human resources, including technical and business skills, and intangible resources, including coordination, change capacity, and risk orientation. Their study demonstrates that AI technology and data alone are insufficient for generating organisational value. Firms also require structures, skills, and an organisational culture that enable value generation from AI investments. Westover (2025) similarly reports that organisational factors, specifically data infrastructure, executive sponsorship, and change readiness, explain 58% of the variance in AI implementation success. Clifton et al. (2020) extend this argument by showing that the consequences of AI are not technologically predetermined, but shaped by how AI is adopted across different economic activities, business cultures, working practices, skill profiles, and places.

AI deployment also has a temporal dimension. Westover (2025), surveying 127 organisations, reports that advanced AI implementation takes on average 2.7 years from initial experimentation to organisation-wide integration, and that technical implementation typically precedes organisational adoption by six to twelve months. This temporal dimension concerns not only the duration of implementation, but also the changing division of labour between humans and AI. Raisch and Krakowski (2021) demonstrate that AI may initially augment human work before automating parts of the same workflow once systems become sufficiently reliable.

Raisch and Krakowski (2021) illustrate this through the example of JP Morgan Chase. In talent acquisition, managers initially worked alongside an AI system to identify reliable predictors of

future job performance. Once the model became sufficiently reliable, the candidate assessment task itself was automated. Human managers nevertheless remained responsible for the subsequent selection stage, where judgement about ambiguous factors such as organisational cultural fit could not be codified algorithmically. The case demonstrates that firms can move through distinct deployment phases, each associated with different implications for automated, retained, and newly emerging tasks. Krakowski et al. (2023) point to a related mechanism, showing that AI may erode the value of existing human capabilities while simultaneously creating new capabilities at the human-AI intersection.

These dynamics do not unfold uniformly across firms. Industry context, task composition, and skill profiles shape how AI affects work. AI adoption progresses more rapidly in knowledge-intensive sectors than in highly regulated industries (Clifton et al., 2020; Westover, 2025). Tasks whose core bottleneck is prediction are generally more amenable to automation than tasks requiring contextual judgement and interpretation (Agrawal et al., 2019). Employees in repetitive and less complex roles are more likely to perceive redundancy risks than those in positions combining technical and interpersonal capabilities (Bankins et al., 2024).

Taken together, these arguments position AI deployment as a firm-level process. AI becomes organisationally consequential through firm-level capability and unfolds over time through implementation and workflow adaptation. The resulting workforce effects vary across industries, task compositions, skill profiles, and the depth of organisational integration. They are therefore firm-specific, rather than technologically predetermined.

2.2. AI Maturity as an Analytical Framework

These firm-level differences create a challenge for comparative analysis. Firms cannot be compared simply by whether they use AI, because similar technologies may remain experimental in one firm while becoming coordinated and embedded in routine workflows in another. What therefore needs to be captured is the extent to which AI-related resources and practices have developed into stable organisational use. AI maturity provides this analytical lens. Drawing on the broader maturity-model literature reviewed by Sadiq et al. (2021), maturity models describe staged organisational development, with each stage reflecting a more advanced configuration of capabilities, structures, and practices. Applied to AI, such models trace how organisations move from isolated experimentation towards systematic, enterprise-wide integration. AI maturity differs from AI readiness. Readiness concerns whether the

preconditions for adoption exist before implementation begins, whereas maturity concerns the extent to which AI has diffused into organisational operations over time (Alsheibani et al., 2018).

Three requirements, drawn from the literature, guide the choice of the AI maturity concept for this study. First, it must be multidimensional. Building on the AI capability perspective established in Section 2.1.2, firms generate organisational value from AI when technological assets are combined with human and organisational resources (Mikalef & Gupta, 2021). Reviewing 15 AI maturity models, Sadiq et al. (2021) find that many models focus primarily on technology, remain descriptive rather than prescriptive, and lack empirical validation. Where models do move beyond purely technical concerns, five dimensions recur, particularly in the empirical frameworks developed by Ransbotham et al. (2020), Hansen et al. (2024), and Weill et al. (2024): data infrastructure as the availability, accessibility, and quality of data feeding AI applications; technical capability as the underlying systems and platforms required for AI deployment; human skills as the combination of technical AI expertise with domain-specific knowledge; cross-functional coordination as the organisation of learning and collaboration across business and technical units; and strategic alignment as the embedding of AI into organisational priorities and leadership commitment. Figure 1 summarises this convergence.

Figure 1: Maturity dimension terminology across three frameworks

AI maturity dimensions Framework	Data infrastructure	Technical capability	Human skills	Cross-functional coordination	Strategic alignment	Empirical basis
Ransbotham et al. (2020)	“Basics”	“Basics”	“Talent”	“Organizational learning”	“Corporate strategy”	Survey, N > 3,000
Hansen et al. (2024)	“Data”	“Infrastructure”	“People”	“Culture”	“Strategy”	Three case studies
Weill et al. (2024)	“Making data accessible”	“Scalable architecture”	“Workforce education”	“Ways of working”	“AI in all decisions”	Survey, N = 721 + interviews

Source: Own compilation based on Ransbotham et al. (2020), Hansen et al. (2024), and Weill et al. (2024).

Second, the maturity concept must capture requirements that arise specifically from AI adoption. Pumplun et al. (2019) demonstrate that organisations require more than general digital transformation capabilities, because AI adoption depends on usable and protected data, iterative experimentation under uncertainty, and employees capable of connecting technical AI knowledge with firm-specific domain expertise.

Third, the maturity concept must account for variation in organisational development paths without implying a uniform stage progression. Hansen et al. (2024) show that comparable maturity states may result from different combinations of resources, priorities, and contextual conditions. The framework adopted in this study therefore needs to distinguish levels of AI integration without assuming that firms proceed through an identical sequence of capability development.

These requirements guide the operationalisation used in this study. The five dimensions identified above ensure that AI maturity is assessed as a multidimensional and AI-specific organisational construct, while the four-stage enterprise AI maturity framework by Weill et al. (2024) provides a structured basis for comparing the interview contexts. The stages are therefore used to classify the extent of AI integration, without assuming that all firms reach these stages through identical developmental paths.

The framework by Weill et al. (2024) distinguishes four stages of enterprise AI maturity. Stage 1, *Experiment and Prepare* (28% of firms in their sample), describes organisations focused on workforce education, acceptable use policies, and initial data access. Stage 2, *Build Pilots and Capabilities* (34%), covers use-case development, metrics, and early process simplification. Stage 3, *Develop AI Ways of Working* (31%), marks a shift towards enterprise-wide integration, scalable architecture, test-and-learn routines, and process automation across broader organisational workflows. Stage 4, *Become AI Future Ready* (7%), describes organisations in which AI is embedded across decision processes and contributes to new AI-enabled value creation (Weill et al., 2024).

2.3. Workforce Displacement as an Organisational Outcome

2.3.1. Defining Workforce Displacement at the Firm Level

If AI maturity captures how far AI has become embedded in a firm's operations, the next question is how this increasing embeddedness may affect the workforce. Before that question

can be examined, workforce displacement must be defined with sufficient precision for firm-level analysis.

Autor (2015) argues that jobs consist of bundles of tasks and that technology substitutes for some tasks while complementing others. Whether the net effect on labour is positive or negative depends on which tasks are affected and on how demand, supply, and organisational choices respond (Autor, 2015; Acemoglu & Restrepo, 2019). At the firm level, workforce consequences are therefore likely to appear first as changes in task composition before they become visible as role elimination or headcount reduction.

Accordingly, workforce displacement refers here to firm-level changes that arise when AI adoption leads to task reallocation, role redesign, role elimination, or skill obsolescence (Autor, 2015; Acemoglu & Restrepo, 2019). Task reallocation occurs when AI assumes responsibility for specific task components, narrowing existing roles or redistributing work across positions. Where task reallocation accumulates within a role without eliminating it, the role is redesigned, meaning that the configuration of tasks, decision rights, and required competencies inside the role changes. Role elimination occurs when positions are removed or phased out because their tasks have been automated or absorbed into other roles. Headcount reduction is treated as a further employment-level outcome that may follow role elimination, for example through layoffs, non-replacement, natural fluctuation, or reduced hiring, but it is not assumed to occur whenever tasks or roles are reconfigured. Skill obsolescence occurs when changes in role requirements reduce the fit between incumbent competencies and the capabilities the redesigned role demands.

AI adoption broadens the forms of skill obsolescence at stake. Earlier technologies primarily affected procedural skills, understood here as the ability to execute defined sequences or process standardised inputs (Autor, 2015). Because AI systems can perform prediction (Agrawal et al., 2019), pattern recognition (Felten et al., 2021), and knowledge synthesis (Eloundou et al., 2024), skill obsolescence may also extend to cognitive forms of expertise. This study therefore distinguishes procedural obsolescence from knowledge obsolescence, where accumulated domain expertise or professional judgement loses part of its relative advantage because AI performs cognitive functions on which that expertise previously depended. In practice, these forms may overlap and may emerge gradually, for example through reduced hiring into exposed roles rather than immediate headcount reduction (Massenkoff & McCrory, 2026).

Workforce displacement must be distinguished from two related concepts. Exposure refers to the extent to which the tasks or abilities associated with an occupation overlap with current AI capabilities (Felten et al., 2021; Eloundou et al., 2024). It identifies susceptibility but does not indicate whether a firm acts on that susceptibility. Automation refers to the process through which tasks previously performed by labour are shifted to capital (Acemoglu & Restrepo, 2019). It is a mechanism that may or may not result in workforce displacement, depending on whether its effects are offset through task creation, demand expansion, or complementary work (Autor, 2015). Workforce displacement therefore captures the organisational outcome that may follow when exposure translates into automation and when automation, combined with organisational decisions about workflow redesign, results in task reallocation, role redesign, role elimination, skill obsolescence, or headcount reduction within the firm.

2.3.2. From Exposure Estimates to Observed Workforce Effects

Having defined workforce displacement as a firm-level outcome, the next issue is whether existing measures can capture it. Prior research has refined the measurement of susceptibility, moving from broad automation-risk estimates to AI-specific exposure measures, but it leaves open whether and how such susceptibility becomes consequential within firms.

Frey and Osborne (2017) assign automation probabilities at the level of occupations and conclude that 47% of U.S. employment faces high computerisation risk. Arntz et al. (2017) demonstrate that this substantially overstates the risk. Workers within the same occupation perform different task bundles depending on their employer, education, and job-specific characteristics. In their recalculation, the share of U.S. workers facing high automation risk falls from 38% at the occupation level to 9% once job-level task variation is incorporated (Arntz et al., 2017). Susceptibility cannot be inferred from occupational titles alone, because the same occupation may contain very different configurations of work.

A second development shifts attention from broad automation risk to AI-specific exposure. Felten et al. (2021) measure the extent to which the abilities required by occupations overlap with current AI applications and find that exposure is concentrated less in manual routine occupations than in cognitive and analytical roles, with financial analysis, accounting, and legal services among the most exposed. Eloundou et al. (2024) extend this analysis to LLMs and find that exposure is particularly pronounced in higher-income and more highly educated work, inverting the pattern associated with earlier waves of automation. Together, these studies

indicate that newer forms of AI extend susceptibility to occupations beyond the routine, low-skill profile that characterised earlier debates.

These exposure measures identify where AI capabilities and occupational task profiles overlap, but they do not show whether firms act on this overlap or what organisational consequences follow. Massenkoff and McCrory (2026) show, by combining theoretical exposure scores with data on actual generative AI usage, that deployment in professional settings remains well below theoretical capability. They do not find broad increases in unemployment among workers in exposed occupations. They do, however, report slower hiring among younger workers entering more exposed roles, suggesting that workforce displacement may emerge through reduced entry-level demand rather than through immediate headcount reduction.

Taken together, the literature points to three conclusions. First, displacement pressures are task-based before they become visible at the level of roles or headcount (Autor, 2015). Second, estimates of susceptibility depend heavily on the unit of analysis, with occupation-level measures obscuring substantial within-occupation variation (Arntz et al., 2017). Third, exposure is not equivalent to realised workforce displacement within firms, and a considerable gap persists between what AI can do and what organisations actually implement (Massenkoff & McCrory, 2026).

2.3.3. Variation Across AI Maturity Stages

If exposure does not by itself indicate realised workforce displacement, variation is likely to arise from how firms deploy AI within organisational workflows. AI maturity, as defined in Section 2.2, captures this variation in the extent of organisational AI integration. At earlier maturity stages, where AI remains limited to isolated pilots or narrow use cases, workforce displacement may therefore centre on task reallocation within existing roles. At more advanced stages, where AI becomes embedded across workflows, the scope for role redesign, skill obsolescence, and eventually role elimination or headcount reduction may increase because AI integration enables workflow redesign at a scale that isolated pilots do not permit. This expectation is examined empirically in the study.

2.4. Ethics and Governance of AI-Enabled Workforce Displacement

2.4.1. Ethical Standards for AI-Enabled Workforce Displacement

Once AI-enabled workforce displacement is understood as a firm-level outcome of deployment choices rather than as a direct consequence of technical exposure, it also requires governance. Decisions about task reallocation, role redesign, skill obsolescence, role elimination, and possible headcount reduction are not neutral operational adjustments, because they redistribute burdens and opportunities across the workforce and therefore require justification to those affected. The literature on AI ethics and related work on workforce transitions identifies fairness, accountability, and supportable transitions as three standards for assessing AI-enabled workforce displacement.

Fairness concerns the equitable treatment of individuals and groups affected by AI-enabled decisions, particularly in relation to bias, discrimination, and unequal outcomes (Jobin et al., 2019; Floridi et al., 2018). AI systems may reproduce existing inequalities when trained on incomplete or socially patterned data (Clifton et al., 2020). Applied to workforce displacement, fairness concerns how burdens and opportunities are distributed. Exposure varies across task bundles and occupational groups (Felten et al., 2021; Eloundou et al., 2024), and within the firm, decisions about which tasks to automate, which roles to redesign, and who is offered reskilling concentrate burdens and opportunities unevenly. Fairness therefore concerns whether this concentration is defensible.

Accountability refers to the assignment of responsibility for AI-enabled outcomes and the capacity to trace those outcomes back to identifiable decision-makers (Jobin et al., 2019; Papagiannidis et al., 2025). Raisch and Krakowski (2021) identify an accountability gap specific to human-AI systems. When prediction, classification, or recommendation is distributed across algorithms, managers, and end users, responsibility for resulting employment consequences becomes difficult to locate, a difficulty reinforced by the interdependent character of AI deployment in organisations (Hillebrand et al., 2025). In the context of workforce displacement, accountability concerns whether responsibility for workforce-relevant decisions remains visible and contestable within the firm.

Supportable transitions draw on the broader concept of sustainable careers, which conceives of careers as sequences of work experiences that provide continuity, meaning, and viability over time (Van der Heijden & De Vos, 2015). Bankins et al. (2024) apply this perspective to AI-related change and argue that these transitions should not be assessed only in terms of efficiency or labour reallocation, but in terms of whether workers can remain professionally viable under

changing role requirements, through reskilling, redeployment, participation, and role redesign. Supportable transitions therefore concern whether workers whose tasks are reallocated, whose roles are redesigned or eliminated, or whose skills lose value remain viable as role requirements continue to change.

Fairness, accountability, and supportable transitions define the normative standards against which AI-enabled workforce displacement can be assessed. They do not, however, specify how firms organise responsibility, review, adaptation, and justification in practice. In this study, governance refers to the organisational arrangements through which firms seek to meet the ethical demands generated by AI-enabled workforce displacement. Papagiannidis et al. (2025) distinguish structural, procedural, and relational governance practices. Structural practices allocate roles, responsibilities, and decision rights. Procedural practices organise assessment, monitoring, review, and correction. Relational practices organise communication, coordination, participation, and capability-building. This study treats these practice types as organisational means that can serve different governance functions. Applied to workforce displacement, two functions are separated analytically, mitigation and legitimisation. Mitigation concerns the organisation of adaptation capacity through which firms seek to absorb displacement pressures before they translate into adverse workforce outcomes. Legitimation concerns the justification of workforce-relevant change to those affected. The same practice may contribute to both functions, but the functions are analytically distinct.

2.4.2. Governance as Mitigation

Building on this distinction, the first analytical function is mitigation. For this function, a structural practice could be the establishment of a central AI office, steering committee, or designated use-case owner responsible for workforce-relevant AI decisions. A procedural practice could be a use-case screening process, workforce-impact assessment, or post-deployment monitoring that makes task and role effects reviewable. A relational practice could be AI-literacy training, role-specific reskilling, internal mobility support, or employee involvement in implementation processes (Hansen et al., 2024; Papagiannidis et al., 2025; Pumplun et al., 2019; Bankins et al., 2024; Westover, 2025). To be defensible, these organisational practices must meet the ethical standards defined in Section 2.4.1 by distributing adaptation opportunities fairly, keeping responsibility for workforce consequences identifiable, and supporting viable transitions for affected workers.

These standards cannot be met through isolated adaptation measures. Bankins et al. (2024) show that AI reshapes role requirements and career trajectories in ways that exceed what workers can manage on their own. Hillebrand et al. (2025) reinforce this point at the organisational level, as AI deployment distributes decision-making, control, and agency across multiple actors and systems. Raisch and Krakowski (2021) add that automation and augmentation unfold over time through successive deployment decisions, with task reallocation in one area potentially reshaping later automation possibilities and adjacent work processes. To meet the standards of fairness, accountability, and supportable transitions, mitigation must therefore be organised as a firm-level function. Structural, procedural, and relational practices need to be connected within that function, because isolated measures cannot address displacement pressures that arise across roles, workflows, and decision processes.

The extent to which firms will need to organise mitigation in this connected form is likely to vary with AI maturity. At lower stages, where AI remains concentrated in pilots, local and reactive responses may be sufficient. As AI becomes embedded across workflows, coordination and adaptation demands increase, requiring clearer cross-unit ownership, more standardised adaptation processes, and broader relational support (Hansen et al., 2024; Papagiannidis et al., 2025). At more advanced stages, mitigation may therefore shift from a project-level response to a standing organisational function. This remains an analytical expectation rather than a proven pattern.

2.4.3. Governance as Legitimation

The second analytical function is legitimation, the firm's explanation and justification of workforce-relevant AI change to those affected. A structural practice that serves this function may be assigning responsibility for explanation and response to a specific department or role. A procedural practice could be documenting the criteria used for automation, role redesign, or access to reskilling and providing a channel through which affected workers can raise concerns. A relational practice could be communication or consultation that makes the workforce consequences of AI-enabled change discussable. To be defensible, these practices must meet the ethical standards defined in Section 2.4.1. They must make the allocation of burdens and opportunities explainable, keep responsibility for workforce-relevant decisions visible and contestable, and show how affected workers can remain viable as role requirements continue to change (Jobin et al., 2019; Papagiannidis et al., 2025; Hillebrand et al., 2025; Bankins et al., 2024).

As noted in Section 2.4.1, the same organisational practice may contribute to mitigation and legitimisation. Reskilling illustrates this point. As mitigation, reskilling produces adaptation capacity by helping workers adjust to redesigned roles. As legitimisation, the same practice signals that affected workers are provided with a credible and justifiable transition pathway, rather than being left to absorb role change individually (Papagiannidis et al., 2025; Hillebrand et al., 2025).

The need for explicit legitimisation is likely to increase with AI maturity. At lower stages, where AI use remains confined to pilots or local task support, workforce effects may appear fragmented, incremental, or too uncertain to justify in detail. As AI becomes embedded in broader workflows, workforce-relevant decisions become more visible and consequential. When AI begins to affect role design or access to adaptation opportunities, affected workers may be more likely to ask why particular tasks are automated, where responsibility for the consequences lies, and which transition options are available. At more advanced maturity stages, legitimisation may therefore move from implicit reassurance towards a more explicit governance function. This remains an analytical expectation rather than a proven pattern.

2.5. Integrated Conceptual Framing

The review established AI as a firm-level phenomenon whose workforce consequences depend on organisational choices about deployment (Section 2.1). It conceptualised AI maturity as the variation in how deeply AI has become embedded in operations (Section 2.2). It defined workforce displacement as firm-level changes in task allocation, role design, role elimination, and skill requirements, including possible headcount reduction, while distinguishing between procedural and knowledge obsolescence (Section 2.3). It identified fairness, accountability, and supportable transitions as relevant ethical standards and framed governance as the organisational response through which firms seek to meet these, distinguishing two functions: mitigation and legitimisation. (Section 2.4).

These literatures have not yet been integrated into a firm-level account of how AI maturity, workforce displacement, and governance relate across stages. Maturity research describes how firms build AI capabilities without addressing workforce outcomes. Displacement research identifies susceptibility and exposure without examining how these translate into firm-level consequences at different stages of integration. Governance research articulates principles and

practices without linking them to the specific displacement pressures that trigger governance demands.

This study addresses that gap by treating maturity, displacement, and governance as interdependent and by examining how they relate across stages of AI integration. The resulting conceptual framework positions AI maturity as the conditioning factor, workforce displacement as the primary organisational outcome, and governance, through its mitigation and legitimisation functions, as the organisational response. It specifies expectations about how displacement and governance are likely to differ across AI maturity stages. Figure 2 summarises these expectations across the four stages proposed by Weill et al. (2024).

Figure 2: Analytical expectations across AI maturity stages

Stages Dimensions	Stage 1: Experiment and Prepare	Stage 2: Build Pilots and Capabilities	Stage 3: Develop AI Ways of Working	Stage 4: Become AI Future Ready
AI integration	Isolated experiments	Systematic pilots	Enterprise-wide workflows	AI embedded in all decision-making
Workforce displacement	Minimal, limited to individual use cases	Localised task reallocation; procedural obsolescence	Role redesign; emerging knowledge obsolescence	Role elimination; possible headcount reduction; broader knowledge obsolescence
Mitigation	Ad hoc, informal	Project-level, reactive	Cross-unit, formalised	Standing organisational function
Legitimation	Not yet salient	Limited, implicit	Increasingly explicit	Standing organisational function

Source: Own synthesis based on the literature reviewed in Sections 2.1–2.4.

The empirical part of this study examines how these expected patterns appear in organisations at different maturity stages, which displacement forms emerge in practice, and how firms organise mitigation and legitimisation in response.

3 Methodology

In this chapter, the empirical research design, data collection process, and analytical procedure are presented.

3.1. Research Design

This study examines how organisations at different AI maturity stages experience and govern workforce displacement. The research question concerns a firm-level relationship between AI maturity, displacement, and governance that is shaped by organisational decisions which quantitative measurement cannot fully capture alone. A qualitative theory-building approach is therefore adopted as it is particularly suited to studies that aim to refine or extend existing frameworks rather than test fixed hypotheses (Eisenhardt, 1989).

The analytical logic is abductive. The constructs developed in Chapter 2 provide deductive anchors, while empirical patterns are allowed to refine, extend or qualify these anchors where the observed material diverges from theoretical expectations.

The unit of analysis is the organisation, classified according to its AI maturity stage.

3.2. Data Collection

Data was collected through semi-structured expert interviews. This format combines comparable coverage of pre-defined topics with the flexibility to follow each informant's organisational experience (Rowley, 2012; Barriball & While, 1994).

Informants were purposively selected based on two criteria. First, they needed sufficient seniority to comment on organisational decision-making and governance. Second, they needed recent direct exposure to AI integration either within their own organisation or in other organisations they worked with. Around 70 potential informants were contacted through LinkedIn, e-mail, professional networks, and snowball sampling, of whom thirteen participated. The sample group combined firm-side practitioners with consultants and one academic informant. Firm-side accounts provided organisation-specific insights, while consultant and academic accounts contributed cross-organisational perspectives.

Sample size adequacy was assessed using the information power framework by Malterud et al. (2016). Information power refers to how much study-relevant information a sample carries. Greater information power reduces the number of interviews required and depends on five

factors: study aim, sample specificity, use of established theory, dialogue quality, and analysis strategy. The narrow study aim concerning firm-level AI maturity and its workforce consequences kept sample requirements moderate. Sample specificity was high because informants had direct professional exposure to AI integration. The established theoretical framework developed in Chapter 2 provided interpretive anchors that further reduced sample requirements. Dialogue quality was supported by 45-minute semi-structured interviews with senior informants. Although the cross-case analytical strategy increased sample demand, the four supporting factors together indicated that thirteen interviews provided sufficient information power for the purpose of this study.

A semi-structured interview guide was developed around the three thematic blocks of the conceptual framework: AI maturity, workforce displacement, and governance. These were preceded by brief contextual questions about the informant's professional role and followed by questions regarding future developments and expectations. The complete interview guide is reproduced in Appendix A.

Interviews were conducted between April and early May 2026 via Microsoft Teams and averaged approximately 45 minutes. Before each recording began, informants were informed about the purpose of the study and the anonymised use of interview material, and verbal consent was obtained. Recordings were transcribed automatically using Microsoft Teams' native transcription function and subsequently checked manually against the original recordings. Interviews were conducted in German, the native language of all participants and the researcher. Following Squires' (2009) guidance on translation in qualitative research, interview excerpts presented in Chapter 4 were translated from German into English to preserve participants' intended meaning rather than literal phrasing.

Each firm-side informant was assessed against the five dimensions defined in Section 2.2 (data infrastructure, technical capability, human skills, cross-functional coordination, and strategic alignment) and on this basis assigned an AI maturity stage from Weill et al.'s (2024) four-stage framework. Figure 3 reports the resulting classifications.

Figure 3: AI maturity classification of firm-side informants

Informant	AI maturity stage	Reasoning				
		Data infrastructure	Technical capability	Human skills	Cross-functional coordination	Strategic alignment
8	1	Ad-hoc (foundational data integration initiated)	Ad-hoc (early LLM and Copilot trials; AI in isolated technical areas)	Ad-hoc (occasional workshops)	Limited (no dedicated coordination structures)	Absent (no strategic direction)
1	2	Developing (internal LLM available; limited data integration at pilot level)	Developing (LLMs and Copilot deployed; first agents pre-operational)	Developing (mandatory training programme)	Developing (Champion network and small workforce-focused unit)	Limited (bottom-up agent development no strategic roadmap)
2	2	Developing (internal LLM available; data handling partly manual)	Developing (company-wide Copilot; no agents yet)	Developing (regular learning content)	Developing (function-specific use of standardised prompts; uneven adoption)	Limited (no dedicated AI structures)
4	2	Developing (functional data team; structured inputs feeding agents)	Developing (functional pilots with measured outcomes)	Developing (mandatory training)	Developing (Champion network across business units)	Limited (unstructured AI work; no defined target picture)
9	3	Established (structured data layer across functional use cases)	Established (custom AI agents and analytical tools)	Established (formal upskilling roadmap)	Established (dedicated future-workforce function)	Established (formal frameworks for use-case selection and resource planning)
13	3	Established (dedicated data foundation function)	Established (enterprise-wide AI deployment with technical standards)	Established (mandatory training; cross-functional working groups)	Established (formal central AI office with multiple working groups)	Established (AI directive; performance tracking; dedicated leadership role)

Source: Own classification based on interview data and the AI maturity dimensions defined in Section 2.2.

Figure 4 summarises the focus area and maturity range of each cross-organisational informant.

Figure 4: Focus area and maturity range of cross-organisational informants

Informant	Type	Focus area	AI maturity range
3	Consultant	Tech due-diligence advisor; assesses AI and capability maturity across target companies in transaction contexts	Stages 1 to 4
5	Consultant	Consultant with public-sector client focus; observes how regulatory and works-council constraints slow AI adoption	Primarily, stage 1 and 2
6	Academic	Researcher on AI and workforce; survey-based cross-industry observations	Stages 1 to 4
7	Consultant	Transformation advisor; observes AI-related shifts in hiring and operating models across client engagements	Primarily, stage 2 to 4
10	Consultant	HR-technology advisor; observes client AI adoption against a baseline of HR-systems digitisation	Primarily, stage 1 and 2
11	Consultant	Transformation advisor; cross-firm observations on AI rollouts and enablement patterns	Primarily, stage 2 to 4
12	Consultant	AI governance and regulatory advisor; observes how compliance requirements shape AI adoption	Primarily, stage 3 and 4

Source: Own classification based on interview data.

3.3. Data Analysis

The interviews were analysed through a Gioia-inspired abductive coding procedure (Gioia et al., 2013), whose rigour criteria were addressed by anchoring every code in verbatim transcript evidence, maintaining a traceable audit trail from first-order codes to aggregate dimensions, and validating all AI-generated suggestions against the transcripts. The procedure proceeded in three stages. First-order codes captured how informants described the topics relevant to this study. Second-order themes were developed across interviews to identify recurring patterns. Aggregate dimensions were constructed to merge related second-order themes into higher-level analytical categories.

The thirteen transcripts were analysed in two steps. First, the researcher manually coded four transcripts selected to represent the main informant types and maturity contexts within the sample. This step allowed the researcher to develop interpretive familiarity with the material

and to produce an initial set of first-order codes that served as a reference for the second step. Manual coding was conducted in Microsoft Excel, with each first-order code linked directly to its corresponding verbatim transcript excerpt.


Second, the researcher coded the remaining nine transcripts with AI support. Two large language models, ChatGPT 5.5 and Claude Opus 4.7, were used in parallel. At each Gioia stage, both models received the research questions and a description of the Gioia procedure. At the first-order stage, the models additionally worked with the nine remaining transcripts. From the second-order stage onwards, the analysis covered all thirteen transcripts, with the first-order codes from both coding steps merged and provided to the models alongside the transcripts. At the aggregate-dimension stage, the previously developed second-order themes were also included. Each model produced an initial coding pass. A verification prompt then asked each model to reassess both outputs against the verbatim transcripts and identify any codes or dimensions without direct verbatim support. In the absence of conventional human intercoder reliability, this provided a structured consistency check. A proposed code or dimension was classified as converging if both models anchored it in the same transcript content, even if they used different labels. It was classified as diverging if they anchored it in different content. However, the AI outputs served only as basis for further analysis. At every step, the researcher reviewed all AI-generated outputs, against the verbatim transcripts and the prior results, namely the manual reference set from the first step and the validated outputs from earlier Gioia stages. Codes and dimensions with direct transcript support were retained, modified, or merged, while those without such support were rejected. The final first-order codes, second-order themes, and aggregate dimensions at each stage were therefore determined solely by the researcher.

4 Findings

In this chapter, the empirical findings from the Gioia-inspired analysis of the thirteen expert interviews are presented. The results are structured around the three elements of the conceptual framing developed in Section 2.5, namely AI maturity, workforce displacement, and governance.

Figure 5 summarises the underlying data structure by linking illustrative first-order codes to second-order themes and aggregate dimensions.

Figure 5: Data structure of the empirical analysis



Source: Own illustration based on interview data, following a Gioia-inspired structure.

4.1. AI Maturity as Formal Coordination with Internal Differences

The first aggregate dimension concerned the organisational form of AI implementation.

When describing AI implementation, the accounts often began with examples of concrete AI tools. Informant 1 referred to an “internal ChatGPT equivalent” and Microsoft Copilot licences, while Informant 9 described an internal RAG agent, a retrieval-augmented generation tool, that worked “like an interview guide” and produced an “opportunity brief and a use case roadmap”. At the same time, informants repeatedly linked tool deployment to the organisational conditions under which AI could become usable in work. Informant 5 argued that firms could “deploy as many tools as you want”, but that without a clear target picture employees would not know “why and for what and how to use them”. Informant 13 similarly framed AI as a “highly strategic topic” involving “change and culture”, with technology being “only one part”. Coordination was also part of the organisational conditions. Informant 13 referred to an AI Hub, a central unit for company-wide AI strategy and implementation, located in corporate development rather than IT, while Informant 4 described AI Champions, employees assigned as contact persons and multipliers for AI use, as responsible for identifying use cases, answering questions, and supporting adoption across business areas. Other accounts connected implementation to the organisational foundations that made tools usable in work. Informant 8 referred to “building the model”, feeding it with organisational data, and “connecting it to organisational systems”, while Informant 11 argued that it was “not enough” to provide a tool without making data visible and internal knowledge accessible.

These organisational conditions appeared in different combinations across the observed accounts. Informant 8 described an implementation form still centred on experimentation, with “hardly any automations yet” and “hardly any clearly defined processes involving AI”, while employees were mainly “trying things out”. Informant 1 described broader tool availability and pilot structures, including an “internal ChatGPT equivalent”, Copilot licences, AI Champions, and a first HR agent that was “not yet truly live”. Informant 4 described a different arrangement built around AI Champions, who existed “for the business areas and for the functional units” and were responsible for “identifying AI use cases” and serving as “someone they can turn to”. The same account also referred to a concrete interview-agent use case, while noting that the agent “only generates value if people start to actually use it”. In other accounts, AI implementation was organised through dedicated units and routines rather than additional roles

alone. Informant 9 referred to “several people working full time” on AI-related questions and described a future-workforce function that advised business units, helped “prioritize use cases”, and assessed implementation options. Informant 13 described the AI Hub through a “hub-and-spoke model”, working groups, focal points, and a use-case process in which a use case was first identified and then moved through “proof of concept, MVP” and “scale-up”. Across the accounts, AI implementation ranged from experimentation and generic access to additional multiplier roles and dedicated coordination routines.

Tool availability did not automatically translate into role-specific use. Several informants described a gap between having access to AI tools and knowing how to use them meaningfully in daily work. Informant 11 summarised this problem by stating that it was “not enough to say: here is Claude, let us go”. Informant 1 made a similar point from the employee perspective, describing mandatory Copilot training but criticising repeated generic prompting guidance: “I don’t need to hear how to write a prompt for the fifteenth time”. Instead, the informant asked for “concrete, practical examples” showing how AI could create “real efficiency” in daily work. Informant 6 also stated that there was “no shortage” of general AI training, but that employees needed application-specific support for their own cases. In this sense, implementation was described not only as making tools available, but as converting access into concrete, time-protected, and work-specific use.

AI implementation also varied within organisations. Informant 2 stated that HR Business Partners were “just starting”, whereas recruiting was “clearly further along” and already worked with standard prompts for job postings. Informant 8 drew a similar distinction between HR and other business or technical areas. While HR was still “in an early phase”, the organisation already used “drones and AI-based analysis to inspect track infrastructure”, where video material was evaluated with the help of AI. The same account also described unevenness within HR itself. Operational HR management was still experimenting, whereas the policy function already worked with AI “extensively”. This unevenness did not only appear in fragmented AI implementation contexts. Informant 9 noted that “business units differ in how far they are”, with some “pioneers” and other units “not moving as quickly”. Informant 11 similarly distinguished between a local office that was “a few steps ahead” of the broader organisation. Informant 13, similarly, stated that implementation “differs from area to area”. Overall, AI implementation appeared to be uneven across functions, sites, business units, and employee groups.

Taken together, the accounts described AI implementation as more than tool rollout. It involved strategic direction, coordination roles, enablement, and data or system foundations. Organisations differed in how these elements were combined. Tool access had to be translated into role-specific use. Implementation also remained uneven across functions, sites, business units, and employee groups.

4.2. Workforce Displacement Concentrated at Task Level

The second aggregate dimension concerned workforce displacement. The findings are structured along the displacement forms defined in Section 2.3: task reallocation, role redesign, skill obsolescence, role elimination, and possible headcount reduction. Stage labels in this section refer to the firm-side classifications in Figure 3 and the cross-organisational maturity ranges in Figure 4.

Task reallocation was observed across all maturity stages and in several functions and forms. Informant 7, observing this in Stage 2 and Stage 3 firms, stated that “repetitive tasks are tackled first”, including “creating invoices, writing emails, reviewing documents”. In procurement, Informant 9 (Stage 3) described an autonomous negotiation solution for smaller purchasing items, where suppliers entered information into workflows and “the AI then negotiates for these cases instead of a human”. Similar shifts appeared in HR and recruiting at earlier maturity stages. Informant 4 (Stage 2) described an interview agent that “generates the questions” and “builds the behaviorally anchored rating scales”, while Informant 8 (Stage 1) referred to everyday HR tasks such as writing emails, drafting warning letters, or creating documents from scratch, where AI helped generate an “initial draft”. In analytical work, Informant 11 (cross-organisational) stated that deep research across “350 sources” could now be done “automatically with one click”. The same account also described AI support in preparing analyses for the relevant audience and output format. Across these examples, informants described task components being prepared, generated, negotiated, or structured through AI-supported systems, while employees remained involved in review, contextualisation, and decision-making.

Role redesign was observed less consistently than task reallocation. Firm-side evidence appeared most clearly at Stage 3, while cross-organisational accounts described related role-level shifts in software-related work. In marketing, Informant 9 (Stage 3) stated that the “marketing role definitely” has changed. Tasks that previously involved coordinating

translation agencies or selecting photographers had become less central, while “prompt skills for AI-based image generation” became more important. The same informant also described a broader scope of responsibility, because translations and language adaptations could now be completed “in a few clicks”, allowing “one person” to cover “several markets”. In software-related work, Informant 11 (cross-organisational) described computer engineers as moving “away from the middle and toward the beginning and the end” of the workflow. In this account, coding itself became less central, while understanding product and business requirements at the beginning and judging whether AI-generated code fits the target architecture at the end became more important. Informant 3 (cross-organisational) made a related point for developers, stating that the relevant profile was no longer “the one who writes good code quickly”, but someone who “captures requirements in a very structured way” and is strong in “requirements management”. In these examples, role redesign concerned a changed weighting of responsibilities, scope, and role-relevant capabilities within existing roles.

Skill obsolescence was described mainly as a relative decline in procedural skills rather than as knowledge obsolescence. Skill-related changes were mentioned across all maturity stages, but the clearest accounts of procedural skill obsolescence appeared at Stage 3 firm-side and across cross-organisational accounts. Informant 7 (cross-organisational), observing this in Stage 2 and Stage 3 firms, described a shift in analytical work from information retrieval to verification, stating that it was “no longer about the first step of getting that information”, but about “validating information” and processing it further. Informant 11 (cross-organisational) similarly noted that established execution skills were losing relative importance, as consultants or investment bankers who were “absolute heroes in Excel or PowerPoint” faced a situation in which this skill was “no longer as important”, while coding skills also became less central in software-related work. Informant 9 (Stage 3) described a similar shift in marketing, where “prompt skills for AI-based image generation” became more important, while employees also needed to “recognize good and bad output”. At the same time, several accounts provided only limited support for knowledge obsolescence. Informant 8 (Stage 1) stressed that “critical thinking” and domain expertise remained important, while Informant 6 (cross-organisational) described AI-supported programming as reducing the importance of “standard coding” skills but increasing the need for “guiding and reviewing the AI”. Overall, the accounts pointed less to the disappearance of professional knowledge than to a changed weighting of procedural execution, validation, quality assurance, and judgement.

Role elimination and headcount reduction were less directly observed than task reallocation, role redesign, and skill obsolescence. In firm-side accounts across Stage 1 to Stage 3, no informants reported AI-related role elimination. Informant 4 (Stage 2) stated that the organisation had “not lost any roles” and had “not eliminated any positions”, while Informant 8 (Stage 1) reported, “I do not think people have lost their jobs because of AI”. Headcount reduction was mentioned, but in a financially motivated context rather than as a case of AI-related role elimination. Informant 9 (Stage 3) reported that “eight positions were reduced” in an innovation and adoption unit “for cost reasons”. Informant 1 (Stage 2) described a similarly ambiguous context, as roles were being “eliminated and rebuilt anyway” due to divestiture, while a recruiter had left because of outsourcing fears rather than AI replacement. Cross-organisational accounts described stronger expectations of future headcount reduction, especially through reduced hiring or smaller teams. Informant 7 (cross-organisational), observing hiring patterns at Stage 2 and Stage 3 firms, referred to “fewer junior positions”, but framed this as a “correlation, not a causation”. Informant 5 (cross-organisational), drawing on observations at Stage 1 and Stage 2 firms, described junior positions as “no longer being filled” and expected “fewer consultants” where managers could produce first drafts themselves with AI support. Informant 11 (cross-organisational) described engineering teams becoming smaller over time, with possible reductions of “30% to 70%”. These accounts connected possible headcount reduction to higher individual productivity and reduced hiring, even where the role type itself continued to exist. Overall, AI-related role elimination and headcount reduction were mainly anticipated rather than directly observed, whereas headcount reduction also was described as possible through reduced hiring, smaller teams, or cost-driven reductions that were not always attributable to AI alone.

Across all maturity stages, informants described several constraints that prevented technical possibilities from turning directly into role elimination or headcount reduction. These included data quality, system integration, legal restrictions, works-council processes, reliability problems, security concerns, and tool quality. Data and system constraints were visible where AI use depended on organisational information being available and connected to existing systems. Informant 8 (Stage 1) described the necessary groundwork as “building the model”, feeding it with organisational data, and “connecting it to organisational systems”. Informant 12 (cross-organisational) similarly described fragmented implementations with “one tool here, another tool there”, while “humans still sit in between” at the interfaces. Reliability and tool-

quality constraints were mentioned where AI outputs still required checking or did not meet professional standards. Informant 1 (Stage 2) described an early HR agent that “worked reasonably well” but “did hallucinate at times”, while Informant 5 (cross-organisational) stated that current AI models could not yet “build really great slides” in consulting work. Reliability concerns were also connected to security, transparency, and legal requirements. Informant 6 (cross-organisational) referred to continuing problems of “reliability”, “transparency”, and “traceability”. Informant 12 added the legal dimension by explaining that, under the EU AI Act, the same technical tool could be treated as legally different systems depending on its use case. For example, Copilot used for document creation and Copilot used for CV screening could fall under different legal requirements, even if both relied on the same underlying tool. Works-council processes formed a separate institutional constraint. In the observed Stage 2 and Stage 3 firm-side accounts, this was already integrated into standard review and compliance procedures around AI use. Works councils are employee-representative bodies that must be involved in certain workplace changes. Informant 1 (Stage 2) stated that the Works Council would have to agree before roles could be removed and that the organisation could not simply say, “We don't need this anymore, AI can do it, bye”. Across the accounts, these constraints showed why task-level changes did not automatically translate into clearly observed AI-related role elimination or headcount reduction.

Overall, the accounts described workforce displacement most directly through task reallocation, role redesign, and procedural skill obsolescence. AI-related role elimination was mainly anticipated rather than directly observed, while headcount reduction appeared as a possible but causally ambiguous outcome.

4.3. Governance Visible Mainly as Mitigation

The third aggregate dimension concerned the governance of AI-related workforce displacement. The findings are structured along the governance functions defined in Section 2.4: mitigation and legitimation. Stage labels in this section refer to the firm-side classifications in Figure 3 and the cross-organisational maturity ranges in Figure 4.

Mitigation was described in firm-side accounts at Stage 2 and Stage 3 and appeared in different organisational forms. At Stage 2, informants referred mostly to enablement and local support. Informant 1 (Stage 2) described mandatory Copilot training, voluntary sessions, access to online learning material, and internal presentations during a digital learning week. Informant 2

(Stage 2) reported regular short learning formats on prompting and possible HR use cases, while Informant 4 (Stage 2) described AI Champions as contact persons for questions about Copilot, licences, and possible use cases. At Stage 3, mitigation appeared more formalised and more closely connected to dedicated organisational structures. Informant 13 (Stage 3) described an AI Hub with working groups for People and Governance, where the People Group focused on “employee training, training concepts, communication formats, and events”, while the Governance Working Group focused on “strengthening the responsible use of AI”. The same organisation had introduced mandatory training, company-wide rules for AI use, and technical standards for developers. Informant 9 (Stage 3) described a future-workforce function that advised business units on AI use cases, prioritisation, implementation options, and workforce effects. Their team examined “how much working time can be saved”, including whether this saved time would result in “staff reduction” or instead give employees “time to redesign part of their work”.

Legitimation was observed less often and usually appeared more indirectly. Across the accounts, AI implementation was often framed through efficiency-oriented rationales, including productivity gains, time savings, and process automation (Informants 1, 4, 6, 9, 11, 13). Legitimation was most visible at Stage 3, in accounts describing dedicated AI roles or communication formats. Informant 13 (Stage 3) described communication formats, AI-related internal events, and top-management panels as ways to address employee concerns. There, the organisation explained that “AI does not replace people” and addressed such concerns directly. Informant 9 (Stage 3) also reported that webinars had included questions about whether AI would replace employees or make them better. However, this was described as one element within a broader programme focused mainly on training, skills, and AI use. In the Stage 1 firm-side account, legitimation appeared more as general reassurance than as direct consequence explanation. Informant 8 (Stage 1) described internal communication in which AI was presented as “a supporting tool, not as a replacing tool”, while also noting that more open communication about possible job effects would be difficult from an employer perspective. This indirectness was described more explicitly in cross-organisational accounts. Informant 12 (cross-organisational) stated that most companies did not openly say that automation might make jobs disappear and that such concerns were more often raised by works councils or staff councils than by management itself.

Overall, governance appeared more visibly as mitigation than as legitimization across all observed maturity stages. Mitigation became more formalised from Stage 3 onwards, with training, support, dedicated AI roles, use-case review, and workforce-effect assessments. Legitimation remained less direct across all stages and appeared mainly as reassurance, support framing, or through employee-representative channels rather than explicit explanation of workforce consequences.

5 Discussion

This chapter interprets the empirical patterns in relation to the conceptual framework developed in Section 2.5 and integrates them into a revised framework in Section 5.4.

The findings partly support the stage-sensitive logic of the conceptual framework. Formal AI coordination, the scope of task reallocation, role redesign, procedural skill obsolescence, and the formalisation of mitigation all increased with AI maturity. At the same time, three patterns held across all observed maturity stages. AI capabilities remained unevenly distributed within firms; workforce displacement remained concentrated at the task level, while AI-related role elimination and headcount reduction remained limited, indirect, or anticipated; and legitimization remained less direct than mitigation throughout the sample. Figure 6 summarises these stage-sensitive and stage-qualifying patterns.

Figure 6: Stage-sensitive and stage-qualifying empirical patterns

Framework element	Stage-sensitive pattern	Stage-qualifying pattern
AI maturity	Formal coordination, ownership, enablement, and use-case routines became more developed with increasing AI maturity.	AI integration remained unevenly distributed within firms across functions, sites, units, and employee groups at every observed maturity stage.
Workforce displacement	The scope of task reallocation broadened with AI maturity, and role redesign and procedural skill obsolescence increased with AI maturity.	Workforce displacement remained concentrated at the task level across all observed maturity stages; AI-related role elimination and headcount reduction were limited, indirect, or anticipated.
Governance	Mitigation became more formalised with AI maturity, from enablement and local support to dedicated AI roles, use-case review, and workforce-effect assessments.	Legitimation remained comparatively indirect across all observed maturity stages and appeared mainly as reassurance, support framing, or through employee-representative channels.

Source: Own synthesis based on interview data and the analytical expectations from Section 2.5.

5.1. AI Maturity as a Dominant Organisational Configuration

In Section 2.5, AI maturity was positioned as the firm-level conditioning factor through which the extent of AI integration was expected to shape displacement and governance patterns. The findings support this conditioning role, but they also qualify what an AI maturity stage represents. The organisational forms of AI implementation described in Section 4.1, namely formal AI structures, routines, and integration mechanisms, became more developed with increasing AI maturity, while internal unevenness across functions, sites, and employee groups persisted across the observed maturity stages. Even in the most formalised setting, pioneers and slower units continued to coexist within the same organisation. Firm-level AI maturity should therefore be understood as a dominant organisational configuration. A dominant organisational configuration refers here to the prevailing organisational pattern around AI, while allowing for variation across internal units.

This qualification refines the maturity logic developed in Section 2.2. Mikalef and Gupta (2021) treat AI capability as the orchestration of technological, human, and organisational resources, while Pumplun et al. (2019) argue that AI adoption depends on conditions that go beyond general digital transformation. The findings extend this argument inside the firm: adoption conditions and AI capabilities may remain unevenly developed across internal units even when the firm as a whole shows a more advanced maturity pattern. AI maturity therefore captures the dominant organisational configuration around AI, rather than the complete internal distribution of AI-related conditions and capabilities.

A staged reading of maturity could imply that internal unevenness is mainly an implementation problem that recedes as firms move toward more advanced AI use. Internal unevenness persisted across the observed maturity range, which weakens the expectation that progression alone would resolve it at later stages. Internal heterogeneity therefore appears as a feature that firm-level classifications can conceal rather than as an early-stage deviation.

AI maturity should therefore not be used as a simple firm-level predictor of workforce effects. It shapes where displacement and governance pressures become visible, but it does not determine them uniformly. Displacement operates at the level of specific functions, roles, and employee groups, where a single firm-level maturity label can conceal the internal variation through which workforce effects emerge. A maturity label should therefore be treated as a

condition for possible workforce effects rather than as evidence that those effects are already realised.

5.2. Workforce Displacement as a Stage-Conditioned Conversion Process

In Section 2.5, workforce displacement was expected to vary with AI maturity, moving from local task reallocation toward broader role redesign, skill obsolescence, and possible employment-level effects. The empirical patterns support this expectation only partially. The scope of task reallocation broadened with AI maturity, but role elimination and headcount reduction remained limited, indirect, or anticipated across all observed maturity stages. This distinction separates AI maturity-driven task reallocation from employment-level displacement. Firm-level workforce displacement should therefore be understood as a stage-conditioned conversion process. This conversion progresses from exposure to task reallocation, then through role redesign and skill obsolescence, to role elimination and headcount reduction. AI maturity affects how broadly task reallocation and role redesign spread across organisational functions, whereas organisational, legal, technical, and institutional filters determine whether and to what extent the conversion develops into role elimination and headcount reduction.

This refinement extends the task-based perspective outlined in Section 2.3. Autor (2015) and Acemoglu and Restrepo (2019) developed the task-based view at occupational and sectoral level, arguing that technology substitutes for certain tasks while complementing others. The present study shows that a similar mechanism operates at firm level, where organisational conditions filter how task reallocation translates into role redesign, skill obsolescence, role elimination, and headcount reduction. The findings also align with Raisch and Krakowski's (2021) automation-augmentation paradox, according to which AI and human work co-evolve at the task level before role-level substitution occurs, if it occurs at all. This study adds that the task-based mechanism also operates as an organisational process filtered by firm-specific structures, practices, and constraints. This also links the conversion process to the maturity refinement in Section 5.1. If AI maturity captures a dominant organisational configuration rather than a uniform firm characteristic, comparable AI capabilities may still generate different workforce effects across units of the same firm.

From the patterns observed across the maturity-stage contexts, four categories of conversion filters emerged: organisational, legal, technical, and institutional. These mediated the translation from task reallocation toward role redesign, skill obsolescence, role elimination, and

headcount reduction, while in some cases actively constraining it. Organisational filters emerged where localised pilots and uneven role-specific enablement prevented isolated AI capabilities from becoming embedded in workflow transformation. Legal and institutional filters operated through compliance requirements and works-council involvement, making the translation of AI capabilities into role elimination and headcount reduction more mediated and less direct. Technical filters arose from reliability thresholds and the need for human review at consequential decision points. These filters also formed part of responsible AI deployment, because they limited or prevented premature delegation of consequential decisions. The same conditions that enabled responsible AI use therefore also constrained the translation of AI capabilities into role elimination and headcount reduction.

This interpretation aligns with the argument by Massenkoff and McCrory (2026) that AI deployment in professional settings remains below its theoretical capability. Their finding that displacement may manifest through reduced entry-level demand rather than through observable layoffs receives partial support in the present study, where cross-organisational accounts reported reduced junior hiring while firm-side informants did not report direct AI-related layoffs.

The findings also qualify the use of occupational exposure measures. Felten et al. (2021) and Eloundou et al. (2024) measure exposure rather than realised firm-level displacement. Where such exposure measures are interpreted as indicators of near-term employment effects, they risk overstating realised displacement because they do not model the organisational, legal, technical, and institutional filters identified here. The concept of conversion filters therefore clarifies why technological capability does not automatically translate into role elimination or headcount reduction.

The empirical evidence synthesised in Section 4.2 and Figure 6 indicates that the stage-sensitive pattern was strongest for the broadening of task reallocation with AI maturity, while AI-related role elimination and headcount reduction remained limited, indirect, or anticipated across all observed maturity stages. The framework's underlying stage logic therefore remains plausible, but the observed evidence supports it unevenly. Because in Section 2.5 the clearest expectation of role elimination and broader knowledge obsolescence at the most advanced maturity stage is located, the present evidence is better read as an incomplete test of that endpoint than as a rejection of it. Another qualification concerns knowledge obsolescence at Stage 3. While role

redesign was emerging as expected, knowledge obsolescence was less clearly established than the framework anticipated. The skill-obsolescence evidence sharpened this qualification. Procedural obsolescence was visible where AI reduced the value of routine drafting, searching, and standardised screening work. Knowledge obsolescence was less clearly established because the evidence pointed more often to the relocation of expertise into review, validation, and judgement than to the disappearance of accumulated domain expertise.

This gap between task reallocation and the other forms of workforce displacement could be read as a timing issue, with stronger role elimination and headcount reduction emerging once firms reach more advanced maturity stages. The conversion-filters argument allows for that possibility but adds that later-stage effects cannot be inferred from AI capabilities or AI maturity alone. Within the observed sample, compliance procedures, employee-representative involvement, and reliability thresholds became more formalised with increasing AI maturity. Their connection to responsible AI governance suggests that they should be treated as mediating conditions in the conversion from task reallocation to role elimination and headcount reduction, rather than as temporary obstacles that advancing AI maturity will inevitably dissolve.

The workforce displacement dimension of the framework therefore holds directionally but not synchronously. Across the observed AI maturity stages, AI maturity broadened task reallocation and produced procedural skill obsolescence, while knowledge obsolescence appeared in fragmented form and role elimination and headcount reduction remained indirect, anticipated, or not directly observable. Task reallocation represented the most AI maturity stage-sensitive form of workforce displacement. Role elimination and headcount reduction appeared contingent on whether the filters between capabilities, workflow redesign, and workforce decisions were loosened, bypassed, or actively reconfigured.

5.3. Governance as Mitigation-Legitimation Asymmetry

In Section 2.5, governance was framed as the organisational response to AI-enabled workforce displacement, with mitigation and legitimation distinguished as two analytical functions. The framework expected both functions to become more developed with increasing AI maturity. The empirical findings support this expectation more clearly for mitigation than for legitimation. Governance was organised through more formal structures, enablement mechanisms, compliance procedures, and employee-representative involvement in the more developed AI contexts. These practices strengthened adaptation capacity, but they did not

produce an equally explicit explanation of what AI-enabled change meant for affected roles and employees.

This pattern qualifies the governance logic developed in Section 2.4. Papagiannidis et al. (2025) conceptualised responsible AI governance through structural, procedural, and relational practices that guide, coordinate, and monitor AI use. The empirical material confirms the relevance of these governance practices, while also revealing an imbalance within them. Structural practices clarified ownership of AI initiatives, procedural practices reviewed acceptable use, and relational practices supported enablement and training. These practices primarily governed AI deployment and the adaptation capacity around it. They connected AI deployment less consistently to an explicit explanation of workforce consequences. This study adds that responsible AI governance should be assessed by how far governance practices make workforce consequences explainable, contestable, and justifiable to those affected, rather than by the presence of formal practices alone.

The divide became particularly visible in organisational communication. Employee-facing narratives framed AI mainly as support, productivity improvement, or augmentation. At the same time efficiency was repeatedly described as a central rationale for AI initiatives. The workforce implications of efficiency, however, remained underspecified. Efficiency gains may mean maintaining output with fewer employees or expanding output with the same workforce. These scenarios involve different workforce consequences. A support-oriented communication frame may reduce employees' concerns, but it does not specify which scenario the organisation expects. In the firms observed here, all of which operated under German co-determination law, works councils held statutory consultation rights and can be read as institutional intermediaries in this process. Legitimation was therefore mediated partly through representative structures rather than through direct management communication alone.

The accountability gap identified by Raisch and Krakowski (2021) and further reinforced by Hillebrand et al. (2025) helps interpret this pattern. Existing literature describes this gap as the difficulty of locating responsibility when decisions are distributed across systems, managers, and end users. The present study extends this logic to organisational discourse. Responsibility remained difficult to locate both for AI-supported decisions and for organisational communication concerning the workforce consequences of AI-enabled change.

The three ethical standards introduced in Section 2.4.1 also help specify this asymmetry. Fairness appeared mainly through rules and review around acceptable AI use, especially where AI entered HR-related or legally sensitive contexts. However, the displacement-specific fairness question remained less explicit, namely how the burdens and opportunities of automation, role redesign, and reskilling are allocated across employee groups. Accountability showed a similar pattern. The accounts described ownership, standards, and review procedures for AI use, but responsibility for explaining workforce consequences was only partly specified. Supportable transitions were visible mainly as broad enablement and training, rather than as explicit redeployment or transition pathways for clearly affected workers. The ethical pattern therefore mirrors the mitigation-legitimation split: organisations were more advanced in organising adaptation than explicitly communicating and justifying the distribution of workforce consequences.

An alternative interpretation is that indirect communication regarding workforce consequences reflects responsible caution under uncertainty. Firms may avoid strong claims about role redesign, skill obsolescence, role elimination, or headcount reduction while future effects remain uncertain. This interpretation captures part of the observed pattern, especially because direct evidence of employment effects remained limited. However, it does not fully explain the coexistence of formal mitigation structures, strong efficiency rationales, and communication that rarely specified the workforce implications of efficiency gains. The legitimation gap therefore reflects more than communication style alone. Instead, it points to an asymmetry in governance development. Across the observed AI maturity stages, the governance progression concerned mitigation, whereas legitimation remained comparatively indirect.

5.4. Integrated Framework

The three refinements developed throughout the discussion do not replace the framework's stage logic, but they specify how that logic operates. AI maturity remained relevant, although its effects were mediated through internal unevenness and organisational filters. With increasing AI maturity, the range of functions and workflows in which AI could affect work expanded. However, this expansion did not translate directly into role elimination, headcount reduction, or explicit legitimation. AI maturity therefore increases the scope of potential workforce disruption while leaving its realised form dependent on where AI capabilities become embedded, how task reallocation is converted into role elimination or headcount

reduction, and how governance practices connect mitigation with consequence explanation. Figure 7 integrates the three empirical refinements.

Figure 7: Empirical refinements to the framework logic

Original framework logic	Empirical qualification	Integrated reading
AI maturity conditions displacement and governance	Firm-level maturity captures the dominant configuration while internal variation persists	AI maturity shapes where workforce pressures become visible
Displacement shifts from localised task reallocation toward role redesign, obsolescence, and role elimination	Task reallocation broadens, while role elimination and headcount reduction remain mediated by conversion filters	Displacement depends on the conversion from capability to workforce effects
Governance develops with maturity	Mitigation formalises more clearly than legitimisation	Governance capacity and consequence explanation develop unevenly

Source: Own synthesis based on interview data and the framework developed in Sections 2.2 and 2.5.

This interpretation preserves the framework's stage-sensitive logic while avoiding a deterministic understanding of AI maturity. The discussion therefore shifts the framework from a sequence of expected stage outcomes toward a conditional explanation of how AI maturity shapes the emergence, conversion, and governance of workforce displacement.

6 Conclusion and Outlook

In this chapter, the research questions are answered, the study's theoretical contributions and practical implications are summarised, and its limitations and directions for future research are outlined.

6.1. Concluding Answers to the Research Questions

This study addressed the absence of a firm-level, stage-sensitive explanation of AI-enabled workforce displacement. Existing research has shown where work is exposed to AI, how organisations develop AI capabilities, and which ethical standards should guide responsible AI use. What remained less clear was how these strands connect inside firms, including how AI

maturity shapes displacement patterns and how organisations govern the resulting workforce consequences. The study therefore examined AI maturity as the conditioning factor through which task, role, skill, and governance effects become visible.

The first subordinate research question asked which workforce displacement patterns are associated with different AI maturity stages. The study found that displacement across the observed maturity range appeared primarily as task reallocation. AI mainly affected preparatory, drafting, retrieval, and first-version tasks, while human workers retained review, judgement, contextualisation, exception handling, and final responsibility. Higher AI maturity was associated with broader task reallocation across functions and workflows. It was not associated, within the observed maturity range, with proportionally clearer role elimination or headcount reduction. Procedural skill obsolescence was visible where routine drafting, manual search, translation or basic coding became less central. Knowledge obsolescence was less clearly established, because expertise more often shifted into review, validation, judgement, and problem framing than disappeared outright. Role elimination and headcount reduction therefore remained limited, indirect, mixed, or anticipated rather than dominant observed effects.

The second subordinate research question asked which governance mechanisms firms use to mitigate and legitimate displacement across maturity stages. The study found that firms governed AI-enabled displacement mainly through mitigation practices, while legitimisation practices remained less developed. Mitigation became more formalised with maturity through AI hubs, working groups, enablement programmes, mandatory training, compliance procedures, use-case review, and employee-representative involvement. These practices helped organisations structure AI deployment, build adaptation capacity, and prevent premature delegation of consequential decisions to AI systems. Legitimation did not develop in parallel. Organisations tended to communicate AI as support, productivity improvement, or capability-building, while direct explanation of what AI-enabled efficiency meant for affected roles and employees remained comparatively indirect. Fairness was most visible as procedural anti-bias protection, accountability was weaker in relation to explaining workforce consequences, and supportable transitions appeared more clearly as broad capacity-building than as explicit pathways for affected workers.

Bringing both subordinate questions together, the study answers the overarching research question as follows. Organisations at different AI maturity stages experience and govern workforce displacement in stage-sensitive but uneven ways. AI maturity broadens the range of functions and workflows in which AI can reallocate tasks and strengthens the formal mitigation infrastructure through which firms manage AI deployment. It does not mechanically produce role elimination, headcount reduction, or explicit legitimization of workforce consequences. AI maturity therefore conditions which forms of workforce disruption can emerge inside a firm. It does not by itself determine whether task-level change converts into role redesign, role elimination, or headcount reduction. Nor does it determine whether firms explain and justify these workforce consequences explicitly. The central conclusion is that AI maturity matters as a condition shaping displacement pressures, not as a direct predictor of realised displacement.

6.2. Theoretical Contributions and Practical Implications

The study contributes to three bodies of literature. First, it contributes to AI maturity research by clarifying how maturity stages should be interpreted in workforce analysis. Existing maturity models capture how firms develop data, technology, skills, integration, and strategic alignment. This study adds that a firm-level maturity classification captures the dominant organisational configuration around AI rather than the full internal distribution of AI capability. This explains why maturity can shape workforce pressures without determining them uniformly across functions, sites, or employee groups.

Second, the study contributes to workforce displacement research by shifting attention from exposure to conversion. Occupational exposure measures identify where AI capabilities overlap with task profiles, but they do not show whether firms translate that overlap into task reallocation, role redesign, role elimination, skill obsolescence, or headcount reduction. The findings showed that task reallocation can be widespread while role elimination and headcount reduction depend on organisational, legal, technical, and institutional conditions. The concept of a stage-conditioned conversion process therefore specifies a firm-level mechanism in the task-based displacement literature. It explains why AI capability may be consequential at the task level without immediately becoming employment-level displacement.

Third, the study contributes to AI governance research by distinguishing mitigation from legitimization. Responsible AI governance is often examined through structural, procedural, and relational practices. The analysis here shows that these practices may strengthen adaptation

capacity without making workforce consequences explainable, contestable, and justifiable to those affected. This distinction shows that the formalisation of AI governance practices and the explicit explanation of workforce consequences can develop separately. A firm may govern AI use increasingly formally while leaving the workforce meaning of AI-enabled change only partly articulated.

In practical terms, senior management and workforce-planning functions should not assess AI workforce risk only through maturity labels, technology inventories, or occupational exposure estimates. They should examine where AI capability is enacted, which task bundles are being reallocated, which roles are being redesigned, and which skills are losing value. Workforce effects may appear before layoffs, through non-backfilling, reduced junior hiring, role narrowing, increasing review burdens, or a shift from execution tasks toward validation and judgement. Workforce planning should therefore track task reallocation and skill change as systematically as headcount movement.

Teams responsible for AI governance, HR functions, and employee representatives should design governance around both mitigation and legitimation. Mitigation requires ownership of AI use cases, review procedures, compliance checks, role-specific enablement, and reskilling. Legitimation requires a defensible account of what AI-enabled efficiency means for workers. Maintaining output with fewer employees and expanding output with the same workforce imply different consequences. Organisations therefore need to specify which scenario they expect, how affected groups will be supported, who is accountable for role redesign and reskilling decisions, and how concerns can be raised. Governance becomes more credible when formal AI controls are connected to an explicit explanation of workforce consequences.

6.3. Limitations, Future Research, and Closing Reflection

Several limitations should be acknowledged. The qualitative expert-interview design supports analytical rather than statistical generalisation. The maturity-stage assignments were interpretive, based on the Weill et al. (2024) framework and the operational dimensions defined in Section 2.2, which leaves room for researcher judgement. Interpretation and coding were conducted by a single researcher with AI support and therefore did not include conventional inter-coder reliability checks. Manual calibration, two-model omission checking, transcript-level verification, and the Excel audit trail reduced this risk.

Further limitations concern the sample. The study combined firm-side, consultant, and academic perspectives. Firm-side accounts provided stronger evidence for what happened inside specific organisations, while consultant and academic accounts provided broader comparison but may have included second-hand interpretations of client organisations or sectoral patterns. The study could not observe a firm-side Stage 4 case directly, because Stage 4 organisations remain rare. Claims about the most advanced maturity stage therefore remain conceptually informed and supported by cross-organisational expert accounts, instead of being directly observed. The firm-side sample was also restricted to large enterprises, which limits transferability to small and mid-sized firms.

Finally, all firm-side accounts were shaped by the German-speaking institutional context, where works councils, data protection rules, and co-determination influence how AI-related workforce consequences are discussed and governed. This limits transferability to settings with weaker formal worker representation.

Future research should first examine the conversion process longitudinally. A central open question is when and under which conditions task reallocation turns into role redesign, skill obsolescence, role elimination, or headcount reduction over time. Longitudinal multi-case studies could track AI use cases from pilot to embedded workflow and examine whether conversion filters weaken, persist, or change as firms mature. This would help distinguish temporary implementation delays from more durable organisational constraints.

Second, future research should more directly examine Stage 4 organisations and firms that have built operations around AI from the start. The framework developed in this study expects the most advanced maturity stage to show stronger role elimination, broader knowledge obsolescence, and more formal legitimisation. Because such organisations remain rare, purposive sampling of high-maturity firms would address this gap. Comparing such firms with conventional incumbents could show whether the conversion filters identified here remain relevant when AI is embedded more deeply from the start.

Third, larger-scale research should test the refined framework across industries, firm sizes, and institutional settings, as small and mid-sized firms in particular may rely on different mechanisms than those observed at the enterprise scale. Quantitative studies could measure firm-level AI maturity, internal unevenness, displacement forms, conversion filters, mitigation

and legitimization practices in the same design. Comparative institutional research could examine whether works councils and similar representation structures narrow the legitimization gap or merely shift where it becomes visible. These studies would clarify which elements of the framework are sample-specific and which apply across organisational scales, industries, and institutional contexts.

The central insight of this study is that AI maturity does not make workforce displacement inevitable. It changes the conditions under which displacement pressures emerge and must be governed. The central question is therefore where AI has become embedded, how task reallocation reshapes roles and skills, and whether the organisation can explain and justify those changes to the people affected.

References

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31–50. <https://doi.org/10.1257/jep.33.2.31>
- Alsheibani, S., Cheung, Y., & Messom, C. H. (2018). Artificial intelligence adoption: AI-readiness at firm-level. In *PACIS 2018 Proceedings* (Paper 37). Association for Information Systems. <https://aisel.aisnet.org/pacis2018/37>
- Arntz, M., Gregory, T., & Zierahn, U. (2017). Revisiting the risk of automation. *Economics Letters*, 159, 157–160. <https://doi.org/10.1016/j.econlet.2017.07.001>
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- Bankins, S., Jooss, S., Restubog, S. L. D., Marrone, M., Ocampo, A. C., & Shoss, M. (2024). Navigating career stages in the age of artificial intelligence: A systematic interdisciplinary review and agenda for future research. *Journal of Vocational Behavior*, 153, Article 104011. <https://doi.org/10.1016/j.jvb.2024.104011>
- Barriball, K. L., & While, A. (1994). Collecting data using a semi-structured interview: A discussion paper. *Journal of Advanced Nursing*, 19(2), 328–335. <https://doi.org/10.1111/j.1365-2648.1994.tb01088.x>
- Clifton, J., Glasmeier, A., & Gray, M. (2020). When machines think for us: The consequences for work and place. *Cambridge Journal of Regions, Economy and Society*, 13(1), 3–23. <https://doi.org/10.1093/cjres/rsaa004>
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.5465/amr.1989.4308385>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308. <https://doi.org/10.1126/science.adj0998>

- Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>
- Hansen, H. F., Lillesund, E., Mikalef, P., & Altwaijry, N. (2024). Understanding artificial intelligence diffusion through an AI capability maturity model. *Information Systems Frontiers*, 26(6), 2147–2163. <https://doi.org/10.1007/s10796-024-10528-4>
- Hillebrand, L., Raisch, S., & Schad, J. (2025). Managing with artificial intelligence: An integrative framework. *Academy of Management Annals*, 19(1), 343–375. <https://doi.org/10.5465/annals.2022.0072>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44(6), 1425–1452. <https://doi.org/10.1002/smj.3387>
- Malterud, K., Siersma, V. D., & Guassora, A. D. (2016). Sample size in qualitative interview studies: Guided by information power. *Qualitative Health Research*, 26(13), 1753–1760. <https://doi.org/10.1177/1049732315617444>

- Massenkoff, M., & McCrory, P. (2026, March 5). *Labor market impacts of AI: A new measure and early evidence*. Anthropic. <https://www.anthropic.com/research/labor-market-impacts-ai>
- Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, 58(3), Article 103434. <https://doi.org/10.1016/j.im.2021.103434>
- Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *Journal of Strategic Information Systems*, 34(2), Article 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- Pumplun, L., Tauchert, C., & Heidt, M. (2019). A new organizational chassis for artificial intelligence: Exploring organizational readiness factors. In *Proceedings of the 27th European Conference on Information Systems (ECIS)* (Research Paper 106). Association for Information Systems. https://aisel.aisnet.org/ecis2019_rp/106
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chu, M., & LaFountain, B. (2020). *Expanding AI's impact with organizational learning* [Research report]. MIT Sloan Management Review and Boston Consulting Group. <https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/>
- Rowley, J. (2012). Conducting research interviews. *Management Research Review*, 35(3/4), 260–271. <https://doi.org/10.1108/01409171211210154>
- Sadiq, R. B., Safie, N., Abd Rahman, A. H., & Goudarzi, S. (2021). Artificial intelligence maturity model: A systematic literature review. *PeerJ Computer Science*, 7, Article e661. <https://doi.org/10.7717/peerj-cs.661>
- Squires, A. (2009). Methodological challenges in cross-language qualitative research: A research review. *International Journal of Nursing Studies*, 46(2), 277–287. <https://doi.org/10.1016/j.ijnurstu.2008.08.006>

- Van der Heijden, B. I. J. M., & De Vos, A. (2015). Sustainable careers: Introductory chapter. In A. De Vos & B. I. J. M. Van der Heijden (Eds.), *Handbook of research on sustainable careers* (pp. 1–19). Edward Elgar. <https://doi.org/10.4337/9781782547037.00006>
- Weill, P., Woerner, S. L., & Sebastian, I. M. (2024, December 19). *Building enterprise AI maturity* (MIT CISR Research Briefing Vol. XXIV, No. 12). MIT Center for Information Systems Research. https://c isr.mit.edu/publication/2024_1201_EnterpriseAIMaturityModel_WeillWoernerSebastian
- Westover, J. H. (2025). Sustainable AI transformation: A critical framework for organizational resilience and long-term viability. *Sustainability*, 17(21), Article 9822. <https://doi.org/10.3390/su17219822>

Appendix

Appendix A: Interview guide

Interview Guide

AI Maturity, Workforce Change, and Governance

Researcher	Yannick Schulz – Católica Lisbon School of Business & Economics
Thesis	AI Maturity and Workforce Displacement: A Qualitative Comparative Study
Format	Semi-structured interview, approx. 35–45 minutes, via video call
Confidentiality	All responses are anonymised. No company or individual names will appear in the thesis.

A. Your Role and Context

1. Could you briefly describe your role and the extent to which you have insight into the use of AI and related workforce changes in organisations?

B. AI Adoption and Embeddedness

2. How would you describe the current state of AI adoption – in your organisation or in the organisations you work with?

3. What role does AI currently play? Is it primarily about efficiency, innovation, or both?

4. Based on your experience, what distinguishes organisations that are further along in their AI journey from those that are still at an early stage?

C. Workforce Changes

5. Can you describe a specific process or work area that has changed as a result of AI? What did it look like before and what does it look like now?

6. Have there also been changes at the level of roles – such as roles being eliminated, not backfilled, redesigned, or newly created?

7. Which skills have become more important through AI, and which have become less important?

D. Governance and Organisational Response

- 8.** When AI changes tasks, roles, or skill requirements – how do the organisations you know respond? What is actually done in practice?
- 9.** Who takes responsibility for the people side of AI-related changes, and how is this coordinated?
- 10.** How are AI-driven workforce changes explained and communicated to employees? And how do they typically react?
- 11.** Do topics such as fairness, transparency, or ethical considerations play a role in how these changes are managed?

E. Reflection and Outlook

- 12.** Looking back, what has worked well in managing AI and its workforce implications – and what has not?
- 13.** Looking ahead three to five years, what further changes do you expect?
- 14.** Is there anything important about this topic that I have not asked about?