Technische Universität München
TUM School of Management

Managing AI

–

Organizational Strategy and Structure in The Between Times of Artificial Intelligence

David Michael Huber

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung eines

Doktors der Wirtschafts- und Sozialwissenschaften (Dr. rer. pol.)

genehmigten Dissertation.

Vorsitz: Prof. Hana Milanov, Ph.D.

Prüfende der Dissertation:

1. Prof. Dr. Oliver Alexy
2. Prof. Dr. Hanna Hottenrott

Die Dissertation wurde am 28.05.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Management am 15.07.2024 angenommen.

ACKNOWLEDGEMENTS

In memory of my father –	forever my inspiration and moral compass.
Dedicated to the next generation –	especially Carlo, Nicolas, Adrian, and Philippa. May you inherit a world where AI works for you and for the good of society, where it has already helped solve the grand challenges of our time, where automation lets you focus on creative endeavors, and where you always have more exciting things to explore than this dissertation.

I am obliged to the many people who have supported me on this doctoral journey.

First, to my supervisor, Oliver Alexy; primarily for being extremely generous with your time and insights and going far beyond any expectations of what being a doctoral supervisor should entail. But also, for introducing me to the world of organization research, for showing me, in theory and by example, what top-level academic contributions look like and how they are created, for always challenging my work, but for being fair, interested, and understanding of the many different challenges of balancing academia and a growing family.

To my examination committee, Hana Milanov and Hanna Hottenrott; for graciously offering your time – not just in the final stretch of this journey, but throughout. The conversations with you as teachers and colleagues were always inspirational and demonstrate how academic excellence is best paired with personal approachability and kindness.

Next, to all colleagues at ERI, both past and present; for the fun and distractions, for making the trips to Garching not a chore but a treat, for the discussions and lunches, and for lightening the load that this journey brings with it. Especially, to the extended Team-O, for being excellent co-authors and supportive colleagues, and more importantly, friends.

I am grateful for the insights of many fantastic scholars with whom I have had the opportunity to discuss my findings, present to, or learn from – amongst them Karim Lakhani, Hart Posen, Henrich Greve, William Ocasio, Markus Reitzig, Maciej Workiewicz, Kristina McElheran, Kannan Srikanth, Stella Pachidi, Matthew Beane, Robert Seamans, and Julien Jourdan; your work and especially your feedback helped strengthen my interest in technological change, my understanding of how organizations function, and my attention to the responsibility they have in society. I hope to contribute a small piece to the fascinating conversations you shape.

I further thank the coordinators of the Strategic Research Foundation, Giada Di Stefano and Tammy Madsen for selecting me as a Dissertation Research Grant recipient very early on in this journey, and the inspirational class of 2020 scholars for the exchanges and input.

Finally, to my family. Elisa, building our life together and our small family over the last years gave me support and distraction when I needed it most – it continues to be an even more exciting journey than this research. To my mother, Kilian, Konstantin, Stef, Ana, Jana; you make me who I am, shape my curiosity, integrity, and perseverance through yours, and are the foundation to everything I do. Thank you all.

I gratefully acknowledge generous financial support from the German Research Foundation (DFG – Project AL 1252/4-1), the Strategy Research Foundation (SRF), and the Bavarian State Ministry for Digital Affairs (BStMD) in cooperation with the appliedAI initiative, without which this research would not have been possible.

ABSTRACT

The introduction and increasing ubiquity of artificial intelligence presents organizations with a set of opportunities and challenges. AI promises to make existing processes more efficient, support complex analyses in decision-making, and increase creativity – but also challenges routines, established hierarchies, and entire business models. The process of introducing AI therefore places organizations and their leadership at an important strategic junction. They must decide how and where to employ AI, how to coordinate the efforts of various departments, and which structures and skills to build – all while having no clear understanding of the potential rewards of various strategies ex ante. From the perspective of organization theory, this resembles the process of innovative search under fundamental uncertainty. Using this lens to analyze the introduction of AI, in this dissertation I ask four specific questions; (i) What are the responsibilities of strategic leaders in this time of uncertainty? (ii) How do organizations navigate the early process of AI introduction? (iii) Which organizational characteristics aid in achieving broad and deep AI integration? And (iv), how might the roles of humans and AI be distributed in the future process of innovation? The first analysis shows how strategic leaders play a key role in determining the degree of AI use for the long-term future of the organization, building a foundation for success. Based on a qualitative study, the process of introduction is then revealed to require dynamic adjustments between centralized and decentralized structures, which align with the appropriate allocation of power. This serves to streamline initially contradicting goals and ultimately create shared mental representations of AI use across departments and HQ teams. Important to this process are the results from the complementary quantitative analysis, which show that the structures supporting the initial use of AI are different from those driving eventual integration breadth, and different again for integration depth. Especially for the latter, I show how a strong focus on department-level structures and investments in complementary assets are key. Finally, I discuss how AI may not only outperform humans in known search spaces, but with increasing data availability may also be able to replicate the mental processes that humans employ to create breakthrough innovation in previously unknown spaces. Overall, this dissertation therefore demonstrates how organizational leaders and users of AI can and should actively manage the introduction of AI, working to create the strategy and structures that serve the identified purposes – and then continuously engage with the technology in order to ensure the creation of constructive, purposeful, and ethical paths through the AI age.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

Abbreviation	Term
AI	Artificial intelligence
API	Application-program-interface
CAWI	Computer-aided web interviews
CDO	Chief digital officer, usually board of management position
CEO	Chief executive officer, usually director of board of management
CFO	Chief financial officer, usually board of management position
CIO	Chief information officer, usually board of management position
COO	Chief operations officer, usually board of management position
CTO	Chief technology officer, usually board of management position
CV	Curriculum vitae
DAO	Decentralized automous organizations
DRUID	Danish Research Unit for Industrial Dynamics
DV	Dependent variable
e.g.	exempli gratia/ example given
et al	and other authors
GDPR	General data protection regulation
GLM	General linear model
GOFAI	Good old-fashioned AI
GPT	General purpose technology
ChatGPT / GPT-3/4	OpenAI's LLM product; Generative Pre-Trained Transformer
HQ	Headquarters of an organization
HR	Human resources
HT	High-tech industry or organization
i.e.	it est/ that is
IMR	Inverse Mills ratio
IP	Intellectual property
IT	Information technology
IV	Independent variable
LLM	Large language model
ML	Machine learning
NA	Not applicable / no value
NGO	Non-governmental organization
NK	Mathematical parameters of a modeled rugged landscape
OECD	Organisation for economic co-operation and development
PCF	Principal component factor
SME	Small- and medium-sized enterprises
SSCI	Social sciences citation index
TRIZ	Theory of inventive problem solving (Original Russian: Teoriya resheniya izobretatelskikh zadatch)
US	United States of America

The Between Times of AI in Organizations

In 1879, Edison famously demonstrated the electric light bulb. Yet twenty years later, only 3 percent of households had electricity. After another two decades, that number accelerated to half the population. For electricity, these forty years were The Between Times. We are now in The Between Times for AI – between the demonstration of the technology's capability and the realization of its promise.

Ajay Agrawal, Joshua Gans, Avi Goldfarb: Power and Prediction (2022b, pp. 3-4)

The reality is that we are already living in the early days of the AI Age, and we need to make some very important decisions about what that actually means. […]
Many people in organizations will play a role in shaping what AI means – but to make those choices matter, serious discussions need to start in many places, and soon. We can't wait for decisions to be made for us, and the world is advancing too fast to remain passive.

Ethan Mollick: Co-Intelligence (2024, pp. 32, 210)

INTRODUCTION

In an era where artificial intelligence (AI) has ceased to be the exclusive muse of poets of the silicon realm and has instead become the diligent apprentice in the workshops of industry, we find ourselves at a crossroads of transformation. The once clear demarcations between human intuition and algorithmic precision are becoming increasingly indistinct, as AI weaves itself into the very fabric of organizational life. This thesis seeks to unravel the complexities of this integration, examining how the sinews of AI not only strengthen but also challenge the skeletal frameworks of contemporary organizations.

At the heart of this exploration lies the question of structural adaptation — how must organizations morph and flex to accommodate this non-human agency that is AI? The inquiry extends into the creation of new teams and departments dedicated to AI, the reengineering of processes to foster AI-human collaboration, the formulation of strategies that leverage AI's predictive prowess, and the recalibration of power dynamics to reflect the new value generated by intelligent systems. Each of these areas is pregnant with potential, yet also fraught with challenges.

The promise of AI in enhancing organizational efficiency is well-documented, with AI's ability to process vast datasets and identify patterns offering a quantum leap in strategic foresight (Davenport & Ronanki, 2018). AI's role in decision-making processes, too, heralds a new dawn of precision and personalization, enabling organizations to tailor their offerings with unprecedented specificity (Brynjolfsson; & McAfee, 2017). Yet, this is not a panacea without its perils. The displacement of jobs by automation remains a specter that haunts the corridors of progress (Acemoglu et al., 2020), and the potential for AI to entrench existing biases — should its training data be flawed — poses a significant risk to equity and fairness (Barocas, Hardt, & Narayanan, 2019).

The ethical dimension of AI's integration into organizational structures cannot be overstated. As AI systems become de facto participants in organizational decision-making, the moral responsibility of their outputs transfers, in part, to the organizations that deploy them. This necessitates a robust ethical framework to govern AI development and deployment, ensuring that these systems do not perpetuate harm and that they operate transparently and accountably (Floridi et al., 2018). The European Union's General Data Protection Regulation (GDPR) has set a precedent in this regard, emphasizing the right to explanation for AI-driven decisions (Goodman & Flaxman, 2017).

Moreover, the power dynamics within organizations are poised for a seismic shift. As AI systems take on more complex tasks, the locus of expertise and authority may shift from human managers to AI systems or those who control them (Lee, 2018). This redistribution of power necessitates a rethinking of leadership models and a recalibration of

```
organizational hierarchies to prevent the emergence of a new digital
divide within the workplace (West, 2018).

In synthesizing these threads, this thesis endeavors to chart a course
for organizations navigating the AI revolution. It is a journey that
requires a confluence of technical acumen and ethical sagacity, a
balancing act between embracing the capabilities of AI and safeguarding
the human essence of organizational life. As we stand on the precipice
of this new era, the insights gleaned from this research aim to serve as
a beacon for organizations striving to harness the power of AI while
maintaining their moral compass.

In sum, the integration of AI into organizational structures is a
multifaceted phenomenon that demands a nuanced understanding of both its
potential and its pitfalls. This thesis is committed to providing that
understanding, offering a comprehensive analysis that will aid
organizations in their quest to evolve alongside AI, ensuring that they
not only survive but thrive in this new landscape shaped by intelligent
machines.
-----------------------------------------------------------------
```

The above passage is the product of a text-based exchange between the (human) author of this dissertation and the open-source large language model (LLM) Llama3 (2024). We see how the model correctly picked up on the challenges highlighted in the text prompt [1], tries to assume the requested personality of a management scholar to further elaborate on their causes and implications, finds suitable academic sources, and attempts to use appropriate vocabulary in conveying its messages. Overall, a very impressive performance – albeit in style perhaps a little too pompous for publication. When work on this dissertation began in early 2020, long-form outputs such as the one above seemed to be more at home in the realm of science fiction than reality, as was the idea of having immediate and constant access to the tools creating

[1] Full text prompt used in the exchange: "You are a senior scholar in the field of management and organization theory. Your research focus is the introduction of artificial intelligence into organizations. Specifically, you are looking at how organizations have to adjust their structures to effectively work with AI; this can include the creation of new teams or departments, new processes, new strategies, the shift of power allocation within the organization. You are well-versed in the vocabulary and style of academic papers in your field. You like to write clearly and concisely. You make sure to consider arguments from a business- and a moral perspective. Please draft an introduction to a thesis about your topic of research, as agreed above. The introduction should be about 500-1000 words in length. It should be structured as follows: Start with a witty and overarching remark on the role of AI in modern society. Then introduce your research focus and the specific areas you are interested in. Briefly summarize the key arguments for each area; how can AI help organizations in this field, what are potential pitfalls, what ethical considerations must organizations take into account. Use citations from renowned institutions or academic publications. As in any academic paper, do not use headers or structure. Write elegant and convincing prose."

them, which, in turn, have constant and live access to large shares of all knowledge created in the history of humankind.

Just 4 years later, we find ourselves in the AI Age (see also Gates, 2023) – but yet, we are still in The Between Times; as the introductory quotes point out, this era is only just beginning and it is safe to assume that current AI tools are the least capable they will ever be (Mollick, 2024). Even though AI already feels omnipresent in our everyday lives, with more and more decisions influenced by algorithmic prediction – from the products we buy, to the routes we drive, the information we consume, and the movies we watch – the extent of our future interactions with the technology, or rather, how much of our lives might remain *un*affected, are uncertain. Governments are struggling to provide political frameworks for future developments as they attempt to balance innovative freedom against personal data protection (Satariano & Kang, 2023), the distribution of wealth in the new world order becomes an increasingly complex problem (Georgieva, 2024), and even the most advanced gurus of the tech-world argue over whether AI promises bring utopia or apocalypse to the human race (Shepardson, Warburton, & Stone, 2023).

The resulting mix of excitement and fundamental uncertainty does not only affect individuals but also envelops organizations, as they, too, attempt to navigate the introduction of this revolutionary technology. The term organization here entails all groups of individuals with identifiable boundaries, created in the pursuit of a common goal – where actors must find ways to manage the division of labor and integration of effort (Puranam, Alexy, & Reitzig, 2014). Large or small, driven by profits or social incentives, whether to score goals, build computers, send satellites into orbit, or perform surgery, organizations are the essential means through which humans create coordinated action (March & Simon, 1958). In trying to make use of AI, organizations are now confronted with a series of challenges. As predicted by Llama3 above, they must decide how and where to employ AI, consider entirely different inputs for decision-making, create structures, strategies, and processes to benefit from efficiency gains, coordinate the efforts of various stakeholders, and manage power allocations – all the while also remain aware of the ethical consequences of AI use. The question of how organizations may navigate these challenges is the subject of this dissertation.

Although phenomenologically acutely interesting, the emergence of AI and the resulting challenges are not, on the surface, necessarily unique from the perspective of organization theory. Especially technological change through IT systems has been a topic of fascination for researchers of organizations for decades, where previous findings have considered such questions as how structures shape the process of technology introduction (e.g., Siggelkow & Levinthal, 2003), how technology affects productivity (e.g., Brynjolfsson, 1993), or how routines of work adapt to fit the technology (e.g., Barley, 1986). In these considerations, the perspective of The Behavioral Theory of the Firm has been particularly useful, and so I build on these ideas in my analysis of the emergence of AI. Arising out of the seminal works of Herbert Simon (1957), Richard Cyert, and James March (1963), the core concept of the Behavioral Theory is the idea of bounded rationality, which fundamentally challenged the previously widely assumed ideals of perfect information and flawless decision-making. Instead, these authors present a novel perspective which allows them to describe the observable behavior of organizations more accurately, looking to understand the internal processes of the organization on a micro-level. Gavetti and colleagues (2012, p. 3) later describe this as "research that opens up the black box of the firm and accumulates theory and evidence on how a firm behaves as a result of lower-level processes"– that is, an attempt to understand the internal workings of an organization, the dynamics of behaviors and choices, rather than perceiving it at as a fixed entity described wholly by its inputs and outputs.

60 years after their introduction, the emergence of AI relates to, and indeed challenges these concepts of the Behavioral Theory on multiple levels. First, the process of organizational innovation is commonly conceptualized as an activity of search (Cyert & March, 1963). The introduction of AI fits neatly into this metaphor; with organizations having

to identify possible application areas that may lie close to (local search) or far from (distant search) their existing ways of organizing. In Chapter 1, I take on the role of strategic leaders of organizations in a stylized process of search, as they embark on their initial AI journey. I systematically analyze the various facets of organizing and look to identify potential improvements through AI, in terms of process-related efficiency and product-related effectiveness, as well as the required changes to the roles of leaders themselves.

In real life, of course, processes of search are messy. All organizational actors are affected by bounded rationality, limited in the amounts of information they can obtain and process, and equally limited in anticipating potential consequences of their actions (Simon, 1955). As a result, organizations may struggle to create a coherent and consistent picture of how to use AI, which must entail the potentially conflicting interests and perspectives of various actors. This also raises questions of who holds decision-making authority and how power and control are allocated. Where organizations previously relied on individuals holding key pieces of information, thereby influencing the direction and goals of search in a dominant coalition, these positions may in the future be held by strategic recommender algorithms, which could themselves become part of a dominant coalition. We must therefore also consider AI as a trigger for structural uncertainty and change. In Chapter 2, I create a process model of how organizations navigate this fundamental uncertainty, based on empirical observations of three multinational corporations and their first steps with AI journey over a number of years. Chapter 3 picks up on these same ideas of the interplay between structure and AI use, and by way of a large-scale survey study, identifies the preconditions that facilitate a successful integration of tools.

While the concepts discussed in Chapters 2 and 3 still fit neatly into the existing framework of agency and control as presented in The Behavioral Theory, AI may also speak directly to the assumption of bounded rationality. This challenges the foundation of The Behavioral Theory and creates a clear conceptual separation between the introduction of artificial intelligence and previous instances of technological change. With ever-increasing computing power, processing speed, and available data to provide a better understanding of the internal and external environments, organizations could conceivably move past the limitations of human rationality in decision-making. The process of problemistic search, as described in the Behavioral Theory, limited by an understanding of the environment and data processing capacities, could thus shift from myopic and initially local to being fully informed and immediately distant. I discuss this opportunity in Chapter 4 of this dissertation, using the example of innovation as the key process of strategic foresight.

Combined, the projects therefore form a cohesive investigative journey, in which we accompany organizations on their path to AI integration – starting with the overarching questions of leadership around which areas and decisions might be particularly affected, moving into the first steps of organizations introducing use cases and attempting to create a shared understanding through variations of structure, corroborating these preconditions on a larger scale, and finally considering the long-term effects of AI integration on the process of innovation. From each study, I derive individual contributions to theory and practice, which are presented in short in the following section. However, considering the collection of studies as a whole, there is a notable red thread running through all analyses:

Strategy matters. Human actors will continue to play vital roles in organizations in the foreseeable future – and have the potential to shape AI use not only for the good of their individual organizations, but for the good of society. At this crucial point in The Between Times, decision-makers possess agency in determining the way forward. For their organizations, leaders can prepare structures and processes that allow a deeper integration and more efficient use of AI tools (see examples in Chapters 2 and 3). For society, leaders must strike the right balance between exploiting these advantages and ensuring that not just legal regulations, but also moral obligations are upheld. AI tools promise to automate many of the decisions that organizations regularly make (see examples in Chapter 1), thereby opening the door for efficiency gains capitalized through headcount reduction, or manipulating the public into the consumption of below-par products and services (see examples in Chapter 4). Alternatively, decision-makers could chose to use the efficiencies created to allow for broader creative experimentation (Mollick, 2024) and enforce human control over the judgement of algorithmic outputs (Agrawal, Gans, & Goldfarb, 2022b) to ensure ethical decision-making. Tying in to both long-standing and recent calls for an increase in attention on the role of managers and the morality of their actions (Aguinis et al., 2022; Hambrick, 1989; Hillman, 2021; Hosmer, 1995), this dissertation therefore serves as a call to take seriously the responsibility that technological developments have placed in the hands of organizational leaders and all users of artificial intelligence.

CHAPTERS OF THIS DISSERTATION

On the Emergence, Types, and Organizational Effects of AI

The chapter immediately following this introduction presents an overview of the history, current capabilities, and key terminology surrounding AI. I introduce two types of classification systems, which categorize AI tools first by their approach to data and learning, and second by their capabilities in recognizing and interacting with their environment. I then discuss how these capabilities shape the organizations attempting to work with increasingly performant machine learning algorithms and how structures and power allocation of organizations are affected. The goal of this chapter is to create a foundational understanding of core technical and strategic ideas on which the analyses in the subsequent studies build.

Chapter 1: The Impact of Artificial Intelligence on Strategic Leadership

In this study, taking the form of a systematic literature review, I investigate the effects of AI on various facets of organizing. From the perspective of strategic leaders, I attempt to answer the overarching questions of how AI will impact organizations – and how strategic leaders should respond. Basing the analysis on the strategy framework by Donald Hambrick (1989), the study discusses how the introduction of AI presents a key change to the organizational setting, which immediately impacts organizational conduct through strategy, structure, and processes, as well as organizational performance through increases in efficiency and effectiveness. The study points out how strategic leaders must make a key decision in shaping the organization to operate *with* AI, that is, using AI as a tool in the existing business model, or *towards* AI, that is, shaping the organization around AI at the core of future operations or products.

This chapter is co-authored with Oliver Alexy and was originally published as a section in the Handbook of Research on Strategic Leadership in the Fourth Industrial Revolution (Huber & Alexy, 2024. Editors: Zeki Simsek, Ciaran Heavey, Brian C. Fox - in print at the time of writing).

Chapter 2: Organizing for AI – Multiple Goals, Structural dynamics, and the Introduction of a General Purpose Technology

This chapter presents a qualitative study, providing an in-depth analysis of how exactly organizations manage the complex process of AI introduction. Using archival data as well as primary interviews from three multinational corporations, the study tracks organizational developments over period of roughly 10 years, from the first beginnings of modern AI use to attempted large-scale implementation and uses these insights to develop a process model. The model reveals how all organizations, independent of their idiosyncratic starting points, iterate between centralized and decentralized structures of decision-making, and how actors in the organization create shared mental representations out of originally conflicting goals.

This analysis contributes three insights to existing literature. Firstly, the process model provides an in-depth understanding of the early stages of structural dynamics and decision-making in organizations as they grapple with new technology presenting fundamental uncertainty. Secondly, in situations of such fundamental uncertainty, the model shows how the centralization of control may be required as an interim step in the creation of shared mental representations. Where centralization was previously often seen as the final step in organizations attempting to exploit new technology after a common understanding is reached, we show how it can also serve to facilitate the initial creation of this understanding, before ultimate use takes place in decentralized, platform-like structures. Lastly, the study speaks to literature on sensemaking, focusing on the role of power struggles emerging out of and ultimately resolving conflicting goals. The study also addressed the important decisions which practitioners must make; often acting against their likely instincts, managers must allow sufficient experimentation, before introducing central structures and process control. To facilitate broad adoption of AI, however, they must then accept a counterintuitive loss of control, going against their experiences from prior technology introduction.

Earlier versions of this paper were presented at the 2023 Academy of Management Conference (Huber, 2023) and the 2024 Organization Science Winter Conference (Huber & Reetz, 2024). The version shown in this dissertation is co-authored by David K. Reetz.

Chapter 3: Measuring Organizational AI Integration – A Quantification of Structural Determinants for AI Breadth and Depth

Building on the ideas presented in Chapter 2, this chapter presents a survey study to answer the question of which kinds of organizations best manage to integrate AI. The dependent variables used in this analysis are binary AI use, as well as integration breadth (measured as the number of use cases in a specific department), and integration depth (measured as the share of departmental processes supported by AI tool). Introducing the two novel measures of AI integration allows me to go beyond the already well-investigated question of ‘who uses AI’ to also ask ‘who uses AI well’ – thereby strengthening the understanding of the effects of firm design choices on technology success. I first confirm existing knowledge by showing that AI use is directly related to organizational-level factors such as business performance, age, and innovation capabilities. I then show, however, that these drivers do not translate to integration breadth, and that integration depth is yet again largely influenced by different structural choices. This analysis emphasizes the importance of departmental-, rather than organizational-level factors in driving AI integration and shows the significance of managerial choices in the interaction with new technology.

I further attempt to build on the idea of decision-making decentralization as discussed in Chapter 2. For this, I employ an exploratory margins analysis to investigate the effects of increasing decentralization on integration depth, for various organizational sizes and points in time of the AI journey. While this final analysis does not yield statistically significant results, the indicative results point in a similar direction to those of Chapter 2; with the understanding that large and small organizations require different types of interventions at different times, leaders should try to allow for dynamic structures of organizing with clearly demarcated shifts between decentralized exploration and centralized control.

A previous version of this chapter is under review by the Journal of Economics and Management Strategy, at the time of writing. The same previous version is also accepted for presentation at the 2024 DRUID conference (Huber, 2024). This study is single-authored.

Chapter 4: Breakthrough Innovation and the Asymptotic Rationality of Artificial Intelligence

In this final chapter, I address the process of innovation as a key aspect of organizing – and one that may be significantly affected by the introduction of AI. This conceptual chapter focuses on breakthrough innovation, which is commonly perceived as a process of search across known and unknown task environments. Such search promises to be upended by the increasing prowess of AI, capable of enhancing the breadth, depth, and speed with which such environments may be searched. While some scholars suggest that AI may surmount the limits of human cognition, others have previously argued that human intuition is inimitable and that humans may therefore always out-perform AI in breakthrough innovation. In this study, my co-authors and I argue that these two perspectives can be brough together in the concept of asymptotic rationality – showing how AI may easily out-perform humans in known task environments for both incremental and breakthrough innovation, but how human intuition plays a larger role in unknown task environments with limited data. We then show how human agents still draw on codifiable techniques – which we summarize as meta-cognition, micro-experimentation, and enforcing preferences – to make unknown task environments amenable to finding breakthrough innovation. Such codifiable techniques may however also be taught to machines, so that those might play an increasingly important role even in unknown spaces.

Previous versions of this chapter were presented at the 2022 Strategic Management Conference (Huber, Reetz, & Alexy, 2022a), the 2022 Vienna Conference on Strategy, Organizational Design and Innovation (Huber, Reetz, & Alexy, 2022b) and the 2021 Soph.I.A. conference (Huber, Reetz, & Alexy, 2021). This chapter is co-authored by Oliver Alexy and David K. Reetz.

On the Emergence, Types, and Organizational Potential of AI

A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

The idea of intelligent machines, capable of replicating or outperforming human thought is anything but new. The fascination with such devices can be traced back to antiquity, as early as the automata hailed as being prophesizing and truth-saying in ancient Greece and Egypt. It then re-emerges throughout history – with the da Vinci design for mechanical calculators in the 1500s, Hobbes ideas of a thinking machine Leviathan in 1650, or Turing's concepts for universal computing machines in the 1930s. (McCorduck, 2004; Norvik & Russell, 2021). The modern history of the academic field, however, begins with the proposal for a summer research project – the famous 1956 Dartmouth Conference. In this proposal, the initiators around John McCarthy define their research interest as being "to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves" and that "the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving" (McCarthy et al., 1955, pp. 2, 13).

From these all-encompassing origins, four more detailed approaches to defining artificial intelligence have emerged (Norvik & Russell, 2021): AI could mean machines that can…

- *Act humanly*, for example with the goal of passing the Turing test through natural language processing, knowledge representation, automated reasoning, computer vision, and human-imitation robotics
- *Think humanly*, for example through machine representations of cognitive functions
- *Think rationally*, building on representations of logics and knowledge of the world to create certainty or probabilities

- *Act rationally*, the most commonly used definition in modern approaches, creating agents that do the right thing by perceiving the environment, adapting to change, and pursuing goals to achieve the best possible outcome

Achieving such intelligence, even at the level of imitating human skills, requires machines to obtain and manage vast amounts of data. Early AI tools were static, so-called expert systems, that had the goal of providing support for decision-makers by solving one particular, clearly defined challenge (Bostrom, 2016). The first such system was the Logic Theorist, demonstrated by Allen Newell, Herbert A. Simon, and Cliff Shaw at the Dartmouth Conference – now widely considered the first "proof positive a machine could perform tasks heretofore considered intelligent, creative, and uniquely human" (McCorduck, 2004, p. 167).

To allow these and subsequently developed tools to work, data had to be "elicited from human domain experts and painstakingly hand-coded in a formal language" (Bostrom, 2016, p. 8), and then updated and validated every time significant changes in the environment would make the provided output useless in their practical application. While costly and complex to implement, expert systems hold their promise; for any given scenario, they are able to replicate the analysis and prediction of domain experts – holding the necessary knowledge to *think humanly* for a situation, given sufficient data input. Now usually categorized as 'Good Old Fashioned AI' (GOFAI), these tools relied heavily on symbol-based logics, using heuristics to search for solutions to specific problems. After the first period of enthusiasm in research, emerging out of the Dartmouth conference, was halted by a clear lack of computing power, the 1980s saw newer technology that made such GOFAI tools or expert systems viable for the first time, leading to a renewed interest in the practical application of the previously mostly theoretical discussion around AI. Their cumbersome handling and performance drawbacks in complex settings, however, soon led to the second AI winter of the 1990s. (Bostrom, 2016).

The alternative to manually hard-coding data and all imaginable decision rules given these data, is machine learning. Learning, of course, is usually associated with humans – the clearest example being children, who learn all quintessentially human skills through an ongoing process of observation, interaction, feedback, and repetition. Basing machine learning on the learning progress of a child was proposed by Alan Turing, who suggested ways of making a machine respond to rewards and punishments, letting it make its own decisions to achieve any goal (McCorduck, 2004). While Turing, as well as other researchers like Newell and Simon, proposed these ideas in the 1940s and 1950s, it took another 50 years of research and development to provide the platforms capable of sustaining such complex endeavors.

No AI has yet been developed that adequately represents the child-machine Turing had envisioned – a seed AI capable of developing general intelligence that addresses a wide range of tasks similar to human intelligence – but after rapid progress in computing power, hard drive costs, and data availability through the internet, machine learning systems are now capable of learning and improvement without requiring human intervention (Brynjolfsson & McAfee, 2014). Where GOFAI systems originally required manual re-configuration to produce different results, modern algorithms using machine learning constantly re-configure themselves based on new training data and their own previous output. This milestone led to all complex use cases shaping our world today and enabled such breakthroughs as autonomous machines with software-hardware interaction, optimization and recommendation systems, speech recognition, medical discoveries, or large language models (Haenlein & Kaplan, 2019).

CLASSIFICATIONS OF AI CAPABILITIES

Approaches to learning

Having established that the goal of modern artificial intelligence is to not only out-think humans, but also become agentic in a way that is target-oriented and self-adapting through machine learning, we can now consider some of the more technical details of *how* algorithms achieve such lofty goals. Following from the first distinction between static GOFAI versus machine learning, a second useful classification addresses the interaction between the system and the training data, that is, how exactly a machine learns. Three approaches are typically considered in this classification; supervised learning, unsupervised learning, and reinforcement learning (for details, see e.g., Lantz, 2019; Norvik & Russell, 2021).

Supervised learning involves training an AI model using pre-labeled data. The model learns to predict outcomes based on identified features from the training data. This is typically implemented using algorithms such as regression-based decision trees, support vector machines, or nearest-neighbors' analyses – which are able to accurately predict values or create classifications on grouped sets of data. This is relevant in use cases such as e-mail spam detection or document classification. However, these implementations require intense preparatory work to create large amounts of training data before an algorithm can be used in production.

Unsupervised learning, in contrast, involves training an AI model using unlabeled data. This means that the model attempts to identify patterns and structures in the provided data, without knowing what patterns are relevant beforehand. The main advantage of this approach, often implemented using k-means and dimensionality reduction or principal component tools, is its ability to handle complex, unstructured data – therefore requiring less intense preparation and cleansing. Such pattern discovery can be critical for market analyses, clustering customer groups based on previously unknown purchasing decisions and behaviors, for data visualization, or fraud detection, recognizing patterns that may have slipped by human analysts. A downside of this approach is that the results may not always be intuitive to humans and therefore challenging to interpret.

Finally, reinforcement learning involves training an AI model to make a series of decisions, with the model learning from the rewards or penalties it receives. Its main advantage is its ability to handle dynamic, uncertain environments, where not even the rules of behavior or responses may be codified – but only the targeted output or reward function is known. Examples of this are robotic control systems, complex games and simulations, or language models. Typical algorithms used for reinforcement learning systems are evolutionary or genetic algorithms, or deep learning neural networks.

This categorization is often (mis)interpreted as being in order of performance, as supervised learning tools are most often used for basic, well-defined and smaller applications and reinforcement learning algorithms are often used for more prominent systems with higher real-world impact. However, the more useful distinction is to be made via efficiency with respect to the availability of well-defined and structured data; wherever large amounts of labeled data are available, supervised learning algorithms are effective and require fewer resources and calculating power. As such, they are employed for individual tasks even in complex fields like chatbots (Mohamad Suhaili, Salim, & Jambli, 2021) medical drug discovery (Singh et al., 2023), or protein folding challenges (Xu et al., 2020). Similarly, lower-complexity algorithms such as decision trees may be recombined in ensembles with more complex tools such as neural networks to find application in advanced use cases such as robotic controls (Colledanchise & Ogren, 2016).

Recognition of self and environment

A third distinct approach to classifying types of artificial intelligence is through their capabilities in interacting with a changing environment. This distinction builds on earlier discussion around self-awareness in psychology (e.g., Baron-Cohen, 1997; Premack & Woodruff, 1978) and was first put into writing by Arend Hintze (2016), with the following categories: Reactive machines, limited memory AI, theory of mind AI, and self-aware AI.

Reactive machine, as the baseline functionality, describes a system that “performs optimally under the right circumstances, [but] does not retain information about earlier inputs, and is thus without an internal state or representations about their environment” (Schossau &

Hintze, 2023, p. 1). These machines perform a single task well, with prepared input data prepared and provided to them, they do not require any connection with the outside environment in order to perform this task, and do not form memories about their previous findings. The expert systems we described earlier clearly fall into this category, as do chess computers like IBMs famous Deep Blue (Hintze, 2016), but importantly, many applications of complex algorithms such as neural networks may also be described as reactive machines, for example when they are used to solve single shot classification problems (Schossau & Hintze, 2023).

Limited memory AI is the adaptation of reactive machines that allow systems to better engage with complex and dynamic environments. As the name suggests, these machines can make use of memory, that is, past data which “is not provided at each instant from an exterior source, but retained within for future use” (Schossau & Hintze, 2023, p. 1). This type of AI is widely implemented today, whenever algorithms are used not for single impressions, but to analyze data over a longer period of time, as is the case in any evolutionary algorithm or recurrent neural networks (Bohm, Kirkpatrick, & Hintze, 2022) – for example in self-driving cars, where data cannot be analyzed recurrently every moment, but previous findings are essential input for decision-making later on.

In contrast to these established processes, theory of mind AI is only just emerging in practice. In psychology and cognitive science, theory of mind describes that, considering an object of investigation, this “individual imputes mental states to himself and to others, [using these] to make predictions, specifically about the behavior of other organisms” (Premack & Woodruff, 1978, p. 515). While this early research showed promise in investigating chimpanzees’ understanding and responses to human emotional cues as part of problem solving, more recent work comes to the conclusion that “after some 35 years of research on mindreading in animals, there is still nothing resembling a consensus about whether any animal can ascribe any mental state” (Heyes, 2015, p. 313). Instead of non-human biological intelligence, theory of mind has now become a focus in the investigation of artificial intelligence. Having an AI understand the emotions of the human it is interacting with promises interesting consequences. The positive effects of such an understanding are clear, for example, increasing social intelligence for human-robotic interactions could make medical or geriatric care more accessible (Winfield, 2018), or improve the interaction of humans and AI in collaborative team interactions by increasing trust (Williams, Fiore, & Jentsch, 2022). At the same time, negative consequences, particularly around the ethical use of AI, are also easy to foresee – with a theory of mind-capable AI possibly first manipulating and then exploiting emotional cues of humans, to extract private data (Cuzzolin et al., 2020), or be used in chatbot- or sales-interactions.

While most experts still treat theory of mind as a theoretical concept in AI, albeit one to be realistically achieved with modern technology (Schossau & Hintze, 2023), others have already begun testing AI on specifically designed tests with mixed results (Xu et al., 2024), albeit indicating that every new generation of large language models achieves consistently better results than previous generations, now solving up to "75% of tasks correctly, on par with six-year-old children" (Kosinski, 2023, p. 16). While these early results are scrutinized as being unreliable in the face of trivial alterations of the prompts used and can therefore not definitively prove the existence of theory of mind (Sap et al., 2022; Ullman, 2023; Whang, 2023) – they show that the concept is increasingly becoming more of a reality than an abstract idea.

Finally, self-awareness in AI, also known as self-representation, or consciousness, requires reflective assessments "not about the environment or other agents, but having information about your own (Theory of) Mind: *Cogito, ergo sum*" (Schossau & Hintze, 2023, p. 2). Such AI systems would recognize themselves as thinking and possibly feeling entities – either developing real emotions or ascribing emotional reactions to their behaviors. Coupled with sufficient processing power and access to new data via internet connections, self-aware

AI might then soon not just act and think human-like, but eventually approach rationality. By developing a sense of its own flaws, forming conceptualizations about its position in the physical and digital world, systems might no longer be under human control, but could strive for continued self-improvement and, ultimately, the singularity of artificial general intelligence (Bostrom, 2016; Tegmark, 2017).

Mapping classifications

Thes above classifications, the first by approaches to learning and the second by the recognition of self and others, are not directly related to each other. In theory, any learning approach could be coupled with any level of environmental recognition – in practice, however, certain combinations are more likely and more practical than others, and more often associated with certain types of algorithms, shown in Table 1, below. Once again, this association is directly linked to the interaction with and availability of data.

Table 1: AI classification structures and frequently used algorithm types (exemplary)

Algorithm types	**Approaches to Data & Learning** *(example use cases)*			**Recognition of self and environment**			
	Supervised learning	Unsupervised learning	Reinforcement learning	Reactive machines	Limited memory AI	Theory of Mind AI	Self-aware AI
Decision trees	Predictive maintenance	.	.	✓	.	.	.
Support vector machines	Document classification	.	.	✓	.	.	.
Clustering	Anomaly detection	Customer segmentation	.	✓	✓	.	.
Dimensio-nality reduction	Speech recognition	Text mining	.	✓	✓	.	.
Topic modeling	Information recommender	Data visualization	.	✓	✓	.	.
Evolutionary algorithms	Scheduling optimization	Resource allocation optim.	Robotic control systems	.	✓	✓	.
Neural networks	Computer vision	Image augmentation	Large Language Models	✓	✓	✓	.

A decision tree algorithm, for example, is almost exclusively used with supervised learning techniques – where labeled data is provided and the algorithm learns to identify and prioritize features, that is, form decision nodes in the branches of the tree, which allow it to best achieve its desired output. Decision trees, and other similar types of algorithms such as support vector machines, are therefore often used as reactive machines; once trained, the nodes of the tree can be used to repeatedly make spontaneous decisions, which require no interaction with past events or memories. This type of logic-tree learning for reactive machines could be used in simple rule-based games, for example. Other types of supervised learning algorithms can also be used as limited memory AI. In this setting, algorithms are trained, but more data is stored in ongoing decision processes and re-used for subsequent decisions. Neural networks can be used in this manner, for example in self-driving cars, creating limited memory AI with supervised learning.

In theory, supervised algorithms could also possess theory of mind. This would entail large, labeled datasets of specific input types that represent human emotional states – for example in images of facial expressions or typical speech patterns used in a wide array of circumstances. With access to such perfectly codified training data representing all possible combinations of emotions, mental states, language, expressions, and actions, an unimaginably complex decision tree *could* achieve theory of mind through supervised learning. However, labeling and making accessible such training data would require extremely resource-intensive work – and would realistically always be subject to errors and limitations in the codification. As such, supervised learning is an impractical approach to achieving theory of mind for AI.

Unsupervised learning, in turn, is most logically associated with limited memory AI. Creating clusters of datapoints in dimensionality reduction or topic modelling, for example, works best when retaining information on previous findings and applying this to new data. This is how use cases in anomaly detection, quality control, or fraud detection are structured. However, unsupervised learning techniques are still bound by pre-determined rules and guidelines, and so attempting to create an unsupervised learning AI with theory of mind would still run into the same problems of codification as supervised AI.

It therefore follows from these descriptions that the most promising avenue to achieving theory of mind AI is through reinforcement learning. Some of the most interesting use cases of today are based on complex algorithms using reinforcement learning systems, such as robotic controls through evolutionary algorithms, or large language models using neural networks with deep learning structures. While evolutionary algorithms and neural networks can also be used for supervised or unsupervised applications, where they serve more clearly delineated goals as reactive machines, in this application, these algorithms are focused on consistent self-improvement, or learning how to become more efficient in their task. By definition, this requires access to previously used data and moving beyond the classification as reactive machines. Access to nearly unlimited data through improved retrieval methods (Lewis et al., 2020), and the ability to openly reconfigure this data no matter the kind of input request without any outside intervention, is what is seemingly allowing these systems to now being taking the leap from limited memory AI to possessing theory of mind.

AI IN ORGANIZATIONS

Structure

In their book Competing in the Age of AI, Marco Iansiti and Karim Lakhani (2020) discuss how organizations need to adjust their data-related operating models in order to reap the benefits of digitization in general, and AI tools in particular (see Figure 1).

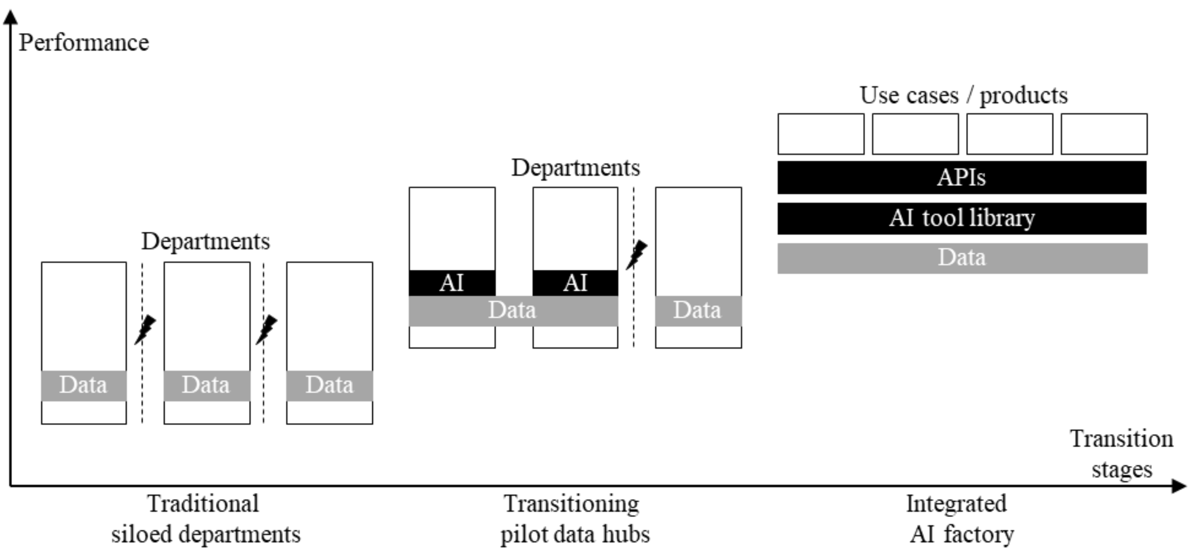

Figure 1: Stages of digital operating model transition (based on Iansiti & Lakhani, 2020, p. 119)

The authors present traditional organizations as characterized by departmental silos. Individual departments may already be working intensely with data and business intelligence tools, but there are strict separations between them that hinder cooperation and alignment. In the initial stages of transformation, some departments may begin data sharing on selected projects, or begin using AI tools, building on such pilot data hubs. However, there is still no overarching AI strategy and departmental silos are still largely in place. In the final stage of transformation, which the authors call the AI factory, departmental boundaries no longer play a role in the organization's approach to data. Instead, the key structures to data operations are use cases; individual data products that span across previous departmental boundaries, built on data lakes that the entire organization has access to through suitable APIs. On top of ensuring data accessibility through aligned formats and hosting services, the data factory also necessitates "powerful algorithms, reusable software components, […] clear policies and governance, dealing with uses from privacy to bias" (Iansiti & Lakhani, 2020, p. 120)

In theory, largely software-driven organizations could choose the integrated AI factory approach not just for their data operating model, but also as the blueprint for their organizational structure as a whole. For established organizations, the data operating model is more likely to be an added layer across existing departmental structures; with new processes and approaches to data handling. In Chapter 1, I reference this as the distinction between working *with* AI as tools versus working *towards* AI as the centerpiece of the organizational strategy (Agrawal, Gans, & Goldfarb, 2022c; Kiron & Schrage, 2019). As Iansiti and Lakhani (2020) make clear, the business model, that is, how to create and capture value, is separate from the data operating model. A successful AI strategy, however, requires intense alignment of the two, regardless of the choice between *with* or *for* AI, to ensure that the data strategy, tool landscape, and creation of accessible APIs take priority over engrained departmental differences – a process of alignment that I analyze more closely in Chapter 2.

Decision-making and power

A transition towards a data-driven operating model should thus be the goal for organizations attempting to work effectively with AI, whether they are aware of it at the outset or not. Managing this transition, however, comes with a set of complicating circumstances, as discussed in the introduction. One of these complications is the potential shift of power that AI may bring in central parts of the organization, which I will discuss in more detail in Chapter 2. This section serves only to provide a first understanding of how AI systems, originally simple tools to support decision-making, can affect organizational dynamics to the extent that they become fundamentally relevant to our conceptual understanding of how organizations function.

In their book Prediction Machines, Ajay Agrawal, Joshua Gans, and Avi Goldfarb (2022c), analyze the component parts of a decision; focusing on the aspects of prediction and judgement. After receiving input data, the focal actor uses their judgement evaluate the data and the potential outcomes – thereby prioritizing what matters to them. This is then used in the prediction component of the decision, where the actor uses their understanding of the relationships at hand to create scenarios associated with certain likelihoods. Combining prediction and judgement leads to action and eventual outcomes of the decision. The authors introduce AI into the framework of decisions as machines that automate prediction – as opposed to judgement, which will remain, for now, in human hands. As AI becomes more readily available, prediction will become more accurate and cheaper. This means that, as its complement in the decision-making process, "judgement, data, and action […] increase in value as prediction becomes cheap. […] In that case, the demand for human judgement will increase" (Agrawal, Gans, & Goldfarb, 2022c, p. 88).

This increase in both the demand and value of human judgement is especially true in light of the complex processes of human behaviors. The prediction of AI tools works best when confronted with well-specified tasks in well-specified environments, where there is sufficient training data available in similar formats and from similar scenarios. This implies, in turn, that AI prediction can fail when considering rare events, or the complex systems of preferences of humans, especially as much of the human “experience is intangible and so cannot be written down or expressed easily” (Agrawal, Gans, & Goldfarb, 2022c, p. 102), taking on the form of latent preferences (March, 1978) . In such situations, with an “absence of good data, our understanding of other humans will provide a role for our judgement skills” (Agrawal, Gans, & Goldfarb, 2022c, p. 110) that are, at the very least, far more difficult for machines to predict with any accuracy.

Finally, this shift can severely impact the dynamics of organizations: “When the implementation of an AI results in decoupling prediction and judgement, there may be an opportunity to increase value creation, but that may require redesigning the system in a way that moves the locus of judgement from the current decision- makers to others. When this happens, power is reallocated. Those who confer judgement ultimately decide and thus have power. New system design that leverages AI may reduce the power of certain individuals who therefore may resist change.” (Agrawal, Gans, & Goldfarb, 2022b, p. 179)

Armed with this background knowledge on AI and its organizational potential, we now turn to the four key questions of this dissertation in the following four chapters.

Chapter 1:
The Impact of Artificial Intelligence on Strategic Leadership

----

Note:
This chapter was previously published as a book section in the Handbook of Research on Strategic Leadership in the Fourth Industrial Revolution (Huber & Alexy, 2024. Editors: Zeki Simsek, Ciaran Heavey, Brian C. Fox - in print at the time of writing).

Artificial intelligence (AI), discussed in technical fields for decades, has gained increasing prominence in the business area since complementary technologies such as cloud storage and multi-core processors have made access to algorithmic tools widely available. As a general purpose technology (Goldfarb, Taska, & Teodoridis, 2023) – pervasive, improving over time, and spawning new complementary technologies (Jovanovic & Rousseau, 2005) – AI is not one single thing, but rather an umbrella term for all artificial system capable of performing actions or producing results previously thought to be accessible only to natural intelligence (Copeland, 2022; Crafts, 2021; Goldfarb, Taska, & Teodoridis, 2023; Iansiti & Lakhani, 2020)[2], a development made clear most recently by the rise of large language processing models (LLMs), ChatGPT and others, stunning the public with their apparent knowledge and skills in a previously unimaginable range of tasks.

To organizations and their leaders, the rise of AI brings a vast range of opportunity, with scholars already describing how it may transform economies (e.g., Autor & Salomons, 2018; Brynjolfsson & McAfee, 2014), open pathways to new business models (e.g., Brynjolfsson; & McAfee, 2017; Iansiti & Lakhani, 2020; Raj & Seamans, 2019), and re-shape

[2] A more specific definition of AI, compatible with this broad understanding, is systems “capable of interacting with the environment by a) gathering information from outside (including from natural language) or from other computer systems; b) interpreting this information, recognizing patterns, inducing rules, or predicting events; c) generating results, answering questions, or giving instructions to other systems; and d) evaluating the results of their actions and improving their decision systems to achieve specific objectives.” (Ferràs-Hernández, 2017, p. 260)

the core organizational activities of task division (e.g., Murray, Rhymer, & Sirmon, 2021) , task allocation (e.g., Tang et al., 2021), information provision (e.g., Waardenburg, Huysman, & Sergeeva, 2022) and reward distribution (e.g., Giermindl et al., 2022).

At the same time, the emergence of AI also brings with it a great number of risks; externally, as it necessitates action against possible emerging competition (e.g., Krakowski, Luger, & Raisch, 2022; Liu et al., 2020), and internally, as it may not just affect individual processes, but then require a fundamental restructuring of entire systems of related processes (e.g., Agrawal, Gans, & Goldfarb, 2022a; Choudhury, 2022), which may lead to the locus of agency and control being shifted in parts from humans to technology (Glikson & Woolley, 2020). Key discussions in the space of AI and management, accordingly, focus on the benefits of augmentation and human-machine interaction on one hand (e.g., Brynjolfsson & McAfee, 2014; Tschang & Almirall, 2021; Verganti, Vendraminelli, & Iansiti, 2020; Wilson & Daugherty, 2018) and the importance of ethics in the use of AI, on the other (e.g., Fjeld et al., 2020; Martin, 2019; Parmar & Freeman, 2016; Wright & Schultz, 2018).

These broad effects of AI on organizations and their environments may render it a strategic inflection point for many organizations, that is, such a foundational change in the technological landscape may lead to “opportunities for strategic leaders to develop new visions, create new strategies, and move their organizations in new directions as they traverse through the turbulence and uncertainty.” (Boal & Hooijberg, 2001, p. 520). Indeed, just how an organization may profit from a general purpose technology such as AI is a vital question for strategy (Gambardella et al., 2021) Given strategic leaders are “the people who have overall responsibility for an organization… [who need] to align the organization with the current and expected external environment… [and] develop an internal organization that has an adaptive capacity and is itself aligned with the strategic thrusts of the firm” (Hambrick, 1989, p. 6), it is their responsibility to address the strategic inflection point of AI emergence –

and to ensure that both themselves and their organizations are adequately prepared to rise to the challenge.

Past research studying at the link between AI and strategic leadership has originated from a variety of different perspectives and disciplines. At a macro-level, authors have for example analyzed the effects of AI on skill composition in labor markets (Tschang & Almirall, 2021), using AI-related job vacancies as an indicator of changing working environments (Acemoglu et al., 2020), and looking at characteristics of roles remaining for human labor (Brynjolfsson & Mitchell, 2017) or the threat of widespread unemployment resulting from automation (Autor & Salomons, 2018). At the firm level, researchers have begun creating a range of detailed how-to guides, describing the steps and preconditions for the successful introduction of AI, ranging from AI as one step in broader efforts towards digitization (Leonardi, 2020), to considerations focusing on the role of data in the AI strategy (Agrawal, Gans, & Goldfarb, 2020; Kruhse-Lehtonen & Hofmann, 2020), to specific guidelines on how to create effective automated decision-support (Watson, 2017), or translate abstract algorithmic results into real-world impact (Davenport & Ronanki, 2018; Grover et al., 2018).

Building on this previous work, in this chapter, we seek to describe how AI may impact strategic leadership and how strategic leadership may enable organizations to benefit from AI by developing and deploying new or improved products and processes. We do so by drawing on the framework presented by Hambrick (1989 - see an adjusted version in Figure 2), in which he defines the management of organizational performance as well as organizational form and conduct as the core responsibility of strategic leaders. Combining this framework with a literature review[3] on the role and effect of AI on strategic leadership and management, we aim to guide strategic leaders towards the key areas of action they must

[3] See Appendix for details on the methodology and identified literature.

address and the personal characteristic they themselves may need to display. To do so, we will highlight how it is particularly important to distinguish the role of strategic leaders in organizations working to make AI a core part of their identity or business models, i.e., *towards* AI, and organizations using AI as a technological tool to enhance existing processes, i.e., *with* AI. At the same time, given the rapid pace at which AI is evolving – at the time of this writing, we have just witnessed the unprecedented explosion of AI fostered by the release of ChatGPT 4 – we expect that our insights may at best capture a lower bound of what AI can and will be able to do. For strategic leaders, that would not only imply that trying to understand *soon* how AI may impact their work will is necessary to avoid losing track of this rapidly evolving field; it also may imply an increasing future shift from AI as a tool to AI becoming a core part of many firm's activities, and, possibly, a veritable competitive threat.

A MODEL OF AI AND STRATEGIC LEADERSHIP

AI and the organizational setting

Before we turn to the specific effects of AI on strategic leadership, we first look at its potential impact on organizations more broadly. This corresponds to Hambrick's perspective on changes in the organizational setting necessitating responses from the strategic leadership. Here, potential changes in the setting comprise both the external environment of the organization, in which AI should play an increasingly significant role, as well as the organization itself, in which other members, possibly without the knowledge of managers, may have already introduced AI to the organization, or may consider doing so.

Machine learning (ML) is the most common approach to implementing AI; broadly defined, ML describes a system through which programs can adjust and improve without requiring any additional input by the programmer. To do this, recursive algorithms use the output of their own calculations and adjust relevant variables in their code structure to better achieve a pre-defined measure of success. While the idea of data-driven, scientific management has been around at least since the Taylor-factories of the late nineteenth century

(Mee, 2022; Wagner-Tsukamoto, 2007), using ML puts data-driven decision-making on steroids: Initially, data tools could only provide *input* to human decision-makers. Process automation tools partly went one step further, as some were able to not just recommend, but make and implement decisions, based on simple if-then-rules. However, adjustments to these rules, as would be common to strategic decisions under uncertainty, were still firmly within the human domain. AI systems, contrary to process automation and related business intelligence tools, may for the first time react independently to changes in the external environment and revise their underlying decision logics. In theory, this opens up the opportunity to remove the human element even from strategic decision-making. That is, as we will further discuss later, AI may even begin to take over certain tasks currently fulfilled by managers themselves, or assist them by taking over crucial activities such as (automated) data analysis and subsequent predictions on the behavior of the external environment, or preparing suitable incentives for employees based on specific personal preferences and experiences.

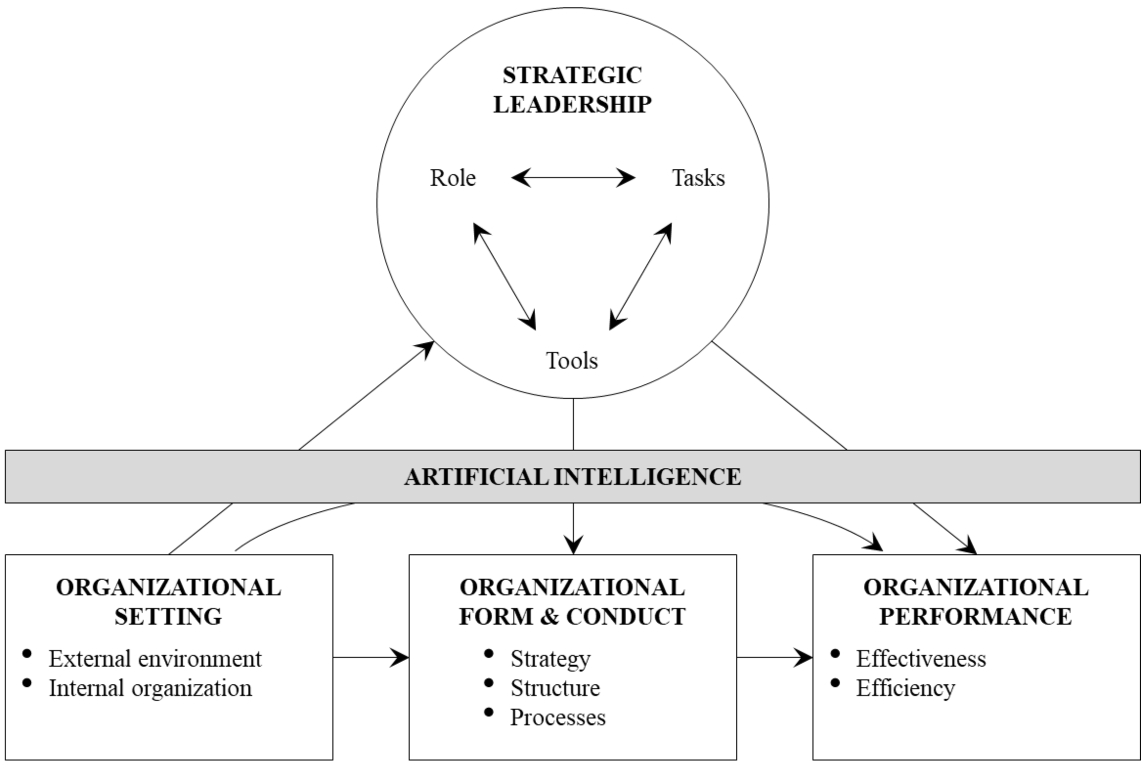

Figure 2: Framework of AI in strategic leadership – adapted from Hambrick (1989)

Nonetheless, barring the existence of a general AI (Bostrom, 2016; Tegmark, 2017), we believe that leaders will still play a vital role in the new world for at least two reasons:

Firstly, current AI systems rely on training data. The quality of an AI system's output is strongly dependent on the amount and quality of data available to feed the algorithm (Sutton & Barto, 2018), which is why AI has only become increasingly prevalent in the age of big data – that is, the availability of a much broader range of codified, machine-readable data on many more aspects of the organizational environment than previously available. They work well in situations where sufficient knowledge on relevant factors exists in codified, readable formats for the algorithm to learn and predict real-world outcomes, and struggle to produce the expected results if the data is flawed or biased (Heaven, 2020). The latest generation of large language models present an interesting mid-point in this transition. These models appear to readily respond to input prompts on nearly any topic, no matter how specific, with the capability of producing output in any required style or complexity of expression. However, they need to heavily rely on the training data available to them. Accordingly, their outputs may be subject to errors based on this data, including a limited understanding of context, the quality of sources, or potential future impact of recommendations (Bogost, 2022; Marr, 2023).

(Good) human managers go beyond this; they may anticipate and react to uncertainties, unpredictable, or unlikely scenarios. Researchers have described these limits of AI in terms of technological barriers to understanding *meaning* (Mitchell, 2019), which limits the applicability of AI on certain tasks (Brynjolfsson & Mitchell, 2017), and therefore warned of unrealistic expectations regarding the general capabilities of AI (Agrawal, Gans, & Goldfarb, 2017) as well as of current implementation projects using AI (Brock & Wangenheim, 2019).

Secondly, even as the technological capabilities of AI will improve over time, encompassing more and more tasks and more and more strategic decisions, we believe humans *should* still be involved and remain capable of interventions. Here, ‘capable’ implies both a process-orientation definition in the sense of having the authority to intervene, and also a skill-oriented definition in the sense of possessing a sufficiently in-depth understanding of the decisions to be made. Especially the second aspect, that human managers must remain informed and acutely aware of the potential errors and dangers of submitting to an algorithm, is already being tested today, in scenarios where ethical concerns are overlooked by machines in many fields (Fjeld et al., 2020).

What, then, does strategic leadership under the influence of AI really mean? As our framework highlights (Figure 2), a significant change in the organizational setting may impact (a) the outputs an organization may produce, (b) the organizational structure it employs to do so, and (c) the guiding and managing role of strategic leaders in this process. Below, we address each of those in turn, looking at how the advent of AI may change existing approaches as well as reveal novel opportunities.

AI and organizational performance

Looking at how AI may affect the bottom line of organizations, Hambrick distinguishes between organizational effectiveness, i.e., the success in achieving an organization’s goals, and organizational efficiency, i.e., the best use of organizational resources required for achieving a certain task. Considering our earlier distinction (working *with* vs. *towards* AI), both types of performance improvements addressed here are more directly related to working *with AI*, meaning AI is used as an automation tool within existing structures, as working *towards* AI would imply changing organizational strategy and conduct (addressed in the next section of this chapter).

The first kind of improvement resulting from starting to work with AI, organizational effectiveness, largely corresponds to thinking of AI as a product technology, that is, it being part of a customer-facing offering of the organization. Famous examples of this are self-driving cars, or recommender algorithms for services such as online streaming or shopping.

Second, AI may increase performance by upgrading existing internal processes, corresponding to improvements in organizational efficiency. This is the case, for example, when algorithms are used to automate HR processes such as the scanning of candidate CVs, or in supply chain management, with the ranking of potential suppliers according to pre-defined criteria, or even in R&D processes, where machine learning in drug discovery, for example, may turn out to be the golden key that will finally help deliver on the promise of high-throughput scanning.

While improvements in each of these areas may take place independently, the two dimensions of product-related effectiveness and process-related efficiency need not be mutually exclusive in practice. One example of a company combining the two is Microsoft, the US-based tech giant. Microsoft has increasingly oriented itself towards AI-based products and services, with the leadership team recognizing a key business opportunity. It became clear that internal operations could also benefit from a more data-driven approach and so thousands of managers and employees started undergoing extensive AI-trainings with the goal of enabling them to recognize the potential benefits of- and successfully introduce AI tools, before also opening up their program to collaborators (Heath, 2018; Roach, 2018).

The examples described above, improvements in organizational efficiency or effectiveness, may be viewed as innovation, as they are the outcome of a (re-)combination of organizational resources to create novel products, services, or processes to reap performance benefits. Accordingly, innovation occurs whenever a new set of resources is being employed or an existing set of resources is employed differently, a definition which encompasses minor adjustments to an existing manufacturing process as much as the creation of entirely new, previously unimagined products. The process of identifying such performance improvements is commonly described as an activity of organizational search (Cyert & March, 1963; Knudsen & Srikanth, 2014; Nickerson & Zenger, 2004). Given human actors are boundedly rational actors in this activity, they are limited in the information they can obtain and process, and in anticipating potential consequences of actions (Simon, 1955). Thus, AI may generally improve firms' capacity to innovate by increasing the power of information processing through machine learning. While this leads to the alluring promise of overcoming the barriers of bounded rationality entirely as we are moving toward general AI, for the foreseeable future, it will be up to strategic leaders to decide on the best use of AI in the context of innovation.

This perspective also shows that, on top of deciding what they hope to gain from innovation (performance gains in efficiency or effectiveness), strategic leaders must set the degree of innovation for which their organization should aim.[4] This choice of scope is commonly described using a logic of incremental versus radical innovation: Incremental innovation implies that organizations try to innovate within search spaces they are familiar with, while radical innovation implies branching out into search spaces unfamiliar to the organization or even creating entirely new ones.

Table 2: Relevance and AI implementation options by innovation types and organizational performance factors

	Organizational performance factors	
Relevant innovation type	**Efficiency** (mainly process-related)	**Effectiveness** (mainly product- & service-related)
Radical innovation (newly created search space)	AI support possible	No AI support possible
Incremental innovation (pre-defined search space)	AI support likely	AI support likely

Jointly looking at (a) the performance factors strategic leaders may focus on and (b) the innovation types they prefer, therefore allows us to systematize how strategic leaders may work with AI. We summarize this in Table 2, to further elaborate, by quadrant, where AI can

[4] Related to this is the choice of how to obtain AI tools; make versus buy. Especially for smaller or resources-constrained organizations, buying ready-to-implement, incremental AI tools is becoming increasingly attractive through cloud services (Watson, 2017). Large enterprises often develop their own tools, to be in control of any intellectual property and gain a competitive edge – which also requires heavy investments into skills and infrastructure and therefore more radical changes to the organization. Indeed, corporations have taken over the forefront of technological development around AI from research institutions (Hartmann & Henkel, 2020).

support increases in organizational efficiency and effectiveness through radical or incremental innovation.

We begin our analysis in the lower-right quadrant of Table 2, where AI should be particularly helpful to organizations hoping to improve their effectiveness incrementally. Given incremental innovation takes place within a pre-defined search space, the possible solutions any agent may find are restricted (Felin, Koenderink, & Krueger, 2017; Helfat, 2021; Rindova & Courtney, 2020). In the context of AI, the search space is represented by available training data – the training data defining the confines within which an AI can operate in the first place. No current AI can detect patterns or move to solutions that lie outside of its determined search space, which would be required for radical innovation to occur. Within a pre-defined search space, however, data on requirements and restrictions may be available and ideally sufficiently codified to function as algorithmic input – requirements for incremental product innovation, often based on large-scale permutations of existing templates to match the required output, such as design or feature combinations. Examples of successes in the area are the much-lauded wins of AI over human champions at complex games such as Go (Chouard, 2016) or Starcraft (Garisto, 2019), or achievements in complex problem solving such as protein folding (Eisenstein, 2021) or compound testing for medical applications (Chan et al., 2019). Organizations will continue achieving successes of this kind driven through AI as a tool for incremental gains in effectiveness.

For radical product- and service innovation (the upper-right quadrant in Table 2), the situation is somewhat different. Radical innovation implies search taking place outside of known problem and solution spaces (Amabile, 2020; Hatchuel, 2002). The ability to enter and interpret such new spaces typically relies on innately human, idiosyncratic abilities, such as the ability to interpret intrinsic values, beliefs, or experiences (Agrawal, Gans, & Goldfarb, 2019; Felin & Zenger, 2009; Gavetti, Levinthal, & Rivkin, 2005; Rindova & Martins, 2017). Whether or not an AI can be helpful in this scenario therefore depends on the amount and quality of available training data – that is, data which precisely maps the complex world of human emotions into machine-readable code. Yet, given the humans in this system may not only struggle to express what kind of innovation they desire (von Hippel, 1994) – meaning no data on their preferences should exist – but also that newly revealed technological possibilities and customer preferences may constantly change the common perception of what a radical innovation should be (Rittel & Webber, 1973), it is unlikely that an AI will be able to collect enough training data to represent the real world adequately (Haefner et al., 2021; Heaven, 2020). We therefore believe that (independently) identifying radically new products and services may continue to lie outside of the scope of AI for the foreseeable future (Amabile, 2020; Huber, Reetz, & Alexy, 2022a).

That said, AI may be seen as enabling human-led radical innovation through an efficiency-lens (upper-left quadrant of Table 2). On one hand, AI may represent a new process tool for developing radical product and service innovation. For example, a first few studies suggest that AI can be drawn on to identify patterns inductively in large amounts of data (Shrestha et al., 2020), which may become be the starting point for humans to conceive of new products and services. This could be, for example, a previously unidentified combination of preferences in food tastes discovered through shopping histories – while an AI may observe these combinations where humans previously had no means to analyze the data, it may still require human understanding to make sense of the new input and conceive of a product combining the tastes. Similarly, AI could eventually outperform humans on the process front if creative meta-techniques, such as design thinking or TRIZ, can be emulated by machine intelligence (Amabile, 2020; Huber, Reetz, & Alexy, 2022a; Verganti, Vendraminelli, & Iansiti, 2020). On the other hand, AI itself may offer previously unconceivable ways of designing and deploying organizational processes, for example in offering access to an amount of data, as a source of process innovation, that would have previously been prohibitively large to analyze. Oftentimes, however, we should be able to treat these, seemingly radical, effects similarly to how AI may lead to increased efficiency through incremental innovation (the lower-left quadrant of Table 2).

Accordingly, the most widely studied effect of AI on organizational performance is through incremental process innovation. One of the most prominent questions in this space is whether AI can and should be introduced as an augmentation- or an automation-technology; that is whether human employees are replaced entirely (automation) or receive selective support from machines in individual tasks (augmentation). Researchers from a range of disciplines have contributed to this debate. From an economic perspective, the impact of automation on jobs and labor displacement is a pressing issue, with authors seemingly converging on the idea that automation is increasing and indeed displacing more labor over time (Autor & Salomons, 2018). Interestingly, while there is a large increase of AI-related job listings since 2010, the direct displacement effects of AI on the labor market at large are not yet noticeable (Acemoglu et al., 2020). Reasons for this could be that effects are difficult to measure due to a lack of precise data (Frank et al., 2019), but also that organizations are expecting results to emerge faster than realistically possible and without the implementation of required complementary technologies (Brynjolfsson, Rock, & Syverson, 2019).

Translated to the level of organizational efficiency, these studies do show that an increasing automation of tasks is a likely scenario for the future. This likely addresses traditional blue-collar jobs, which are at risk of replacement especially when AI is coupled with increasingly precise robotic systems in manufacturing, but may also affect process-driven white-collar roles in areas such as data entry or document management (Brynjolfsson & Mitchell, 2017). Large language models may continue to extend this list of threatened roles, moving, for example, into routine text creation for online articles or social media posts, perfectly crafted to increase audience engagement (Davenport & Mittal, 2022), at a rate unobtainable to human authors. Strategic leaders are therefore put in a position of responsibility on more than one level; not only to achieve the strongest performance benefits for their organization, potentially including strong automation in certain industries, but to also maintain a human-centered perspective of work, ensuring re-skilling opportunities or other compensatory measures for employees threatened by lay-offs.

Even for professions at lower risk of full automation, primarily defined by a larger share of non-repetitive tasks without clearly codifiable inputs or outputs (Tschang & Almirall, 2021), AI may still play a role in the future – in the form of task augmentation or human-AI-collaboration. Researchers have produced a range of models describing how human and AI may work together in various forms (Murray, Rhymer, & Sirmon, 2021). The model created by Athey, Bryan and Gans (2020) focuses on the possible loci of decision-making and shows that the most collaborative augmentation of jobs may be achieved through high-quality AI operating under human decision-making authority. Puranam (2020) instead addresses the interdependence of tasks, showing that the greatest benefit from augmentation may be reaped when there is no clear superiority of performance between human and artificial agents, but the errors of both types of agents are complementary and therefore cancel each other out when combined in a joint decision.

These models focus on decision-making, i.e., complex tasks that require reflection and analysis, which we will discuss in more detail in the next section. A more direct effect of AI on organizational performance, however, comes through the augmentation of jobs through automation of individual sub-tasks that previously required time as a costly human resource, without providing much return for that time. Sub-tasks of this nature are, e.g., document management, the scanning and input of contract or billing data, or the supervision of machinery operations. In automating these sub-tasks, whilst still maintaining the role itself and the host of other sub-tasks it entails, organizations should be able to free up their human resources to focus on more productive tasks – or even allow for more organizational slack, which is seen as a contributor to innovation (Cyert & March, 1963; Nohria & Gulati, 1996). Tang and colleagues (2021) further qualify this argument by showing that employees lower in conscientiousness, although usually seen as an indicator for lower job performance in the modern work environment, might benefit *more* from the support of AI than their colleagues with higher conscientiousness. They explain this by arguing that employees with naturally high conscientiousness already perform well at the tasks to be automated and therefore stand to benefit less, compared to their colleagues who would organically under-perform at those tasks. Other researchers have found that higher-skilled employees benefit *more* from AI tools, especially in creative roles, as they are more accepting to and capable of working with AI-generated suggestions (Jia et al., 2023). Designing roles to fit AI tools to the more suitable kind of employee according to specific goals could therefore be a key qualifier for strategic leaders wanting to benefit directly from the introduction of AI as an efficiency-increasing tool.

In summary, while we believe the effect of AI on improvements in organizational effectiveness through radical innovation to be limited in the short run, algorithmic search can certainly and significantly improve all incremental product and service innovation. The managerial implications of this are clear: Ensuring that (i) machine learning competencies are built up wherever the goal incremental innovation with sufficient training data and (ii) human innovative competencies are maintained in parallel, as this will still be the only way to create truly novel goods and services. At the same time, we urge managers and researchers to maintain an open mind and continue engaging with the possibility that AI systems may support them in experimentation even in new search spaces, thereby preventing a narrow focus on the application of new technologies unintentionally leading to myopia (Balasubramanian, Ye, & Xu, 2020).

In addition, we note how AI has the potential to fundamentally shift our perspective on search within known spaces (i.e., incremental innovation). In this setting, organization theory traditionally distinguishes local and distant search (e.g., Argote & Greve, 2007; Cyert & March, 1963; Puranam et al., 2015). As an AI will be capable of cycling through model-solutions within a known search space with ever-increasing speed regardless of the distance of solutions to the starting point, the distinction between local and distant search may become obsolete. Instead, as the previous paragraphs indicate, we believe that the distinction between incremental and radical innovation will become more relevant.

AI and organizational form and conduct

Beyond changing what organizations do, AI also offers opportunities for strategic leaders to change how they may design and govern the organization through strategy, structure, and processes. How an organization draws on AI to change its form and conduct should ideally correspond to whether it hopes to use AI to improve its effectiveness, efficiency, or both (e.g., Donaldson, 2001). In this analysis, we follow the distinction proposed by Kiron and Schrage (2019), who separate between strategizing for and with AI, and expand their logic to *organizing towards* AI, i.e., creating a setting for AI as a core element of the organization's future, and organizing with AI, i.e., using AI as a tool in the process of creating and organizational form, and extend their distinction to organizational structure and processes.

Table 3: AI relevance in areas of organizational form and conduct (selected examples)

	Goal of AI usage	
Aspect of organizational form/conduct	**Towards AI** – AI as core organizational goal	**With AI** – AI as tool for other organizational pursuits
Strategy	Position AI as key part of future business model, e.g., feature in product or service	Analyze market and competitor data more deeply and efficiently
Structure	Restructure departments to suit AI-based operations	Analyze organizational networks to assess communication streams and proximity of departments
Process	Ensure awareness and usage of AI tools; increase acceptance and trust in results	Automate tasks/roles (see section on organizational efficiency) & supervise human employees

AI and strategizing

Strategizing, by definition, is one of the key tasks of strategic leaders. Hambrick (1989, p. 6) summarizes strategizing as "align[ing] the organization with the current and expected external environment – technology, market trends, regulatory forces, competitor actions, and so on."

To better understand the impact of AI on this extensive process, we follow existing research in separating the *analysis* part of strategy definition, i.e., understanding the competitive situation, from *formulation*, i.e., devising appropriate responses, and *implementation*, i.e., orchestrating the following organizational changes (Schendel & Hofer, (eds.). 1979; Simsek, Heavey, & Fox, 2021).

Strategic competitive analysis is where we expect the strongest impact of AI on the roles of strategic leaders to emerge. Data analysis and managerial intuition have long gone hand-in-hand in the process of strategizing, with Herbert Simon comparing managers to chess grandmasters and stating that "the experienced manager, too, has in his or her memory a large amount of knowledge, gained from training and experience and organized in terms of recognizable chunks and associated information" (Simon, 1987, p. 61).

With the trend of datafication continuing to proliferate, we expect the amount of available and usable data for organizations to increase. This becomes training data for AI, allowing systems to become increasingly performant, delivering reliable results more often. This trend was already recognized decades ago, with Simon noting "the body of evidence from artificial intelligence research indicates that expert computer systems, capable of matching human performance in some limited domain, can be built by storing in computer memory tens of thousands of productions"[5] (Simon, 1987, p. 60). Focusing on purely the analysis of data, we see AI taking on the dominant role in the strategy process, becoming a more reliable analyst than any human expert or manager.

Human intuition, earlier already praised as a necessary ingredient for radical innovation, will still have a role to play in this step, however. The reliability of results

[5] Simon refers to productions, implying simple if-then-commands. This shows two strong developments that have taken place: (i) Modern systems no longer use hard-coded commands, but instead rely on machine learning to continuously improve and (ii) the number of commands to be stored in a machine has increased exponentially, making modern systems more performant than Simon and his peers would likely have imagined in their wildest fantasies.

generated by AI systems will always be contingent on the human experts creating and monitoring them (Kiron & Schrage, 2019). This entails choices to be made; parameters to be defined for the algorithm to work within, and data to be selected to feed to the algorithm for training. Should these factors be askew, so will be the results. Of course, an algorithm can be programmed to self-adjust, should the results delivered not achieve expected levels, but the definition of what these results *should* look like in algorithmic terms is up to the developers creating the algorithm. Balasubramanian, Ye and Xu (2020) warn that the use of AI might, under certain conditions, even increase organizational myopia and thereby negate all potentially positive effects of faster and more reliable data analysis.

The larger question at hand is therefore; under what circumstances and for which components of strategic decision-making will AI improve results? In their model of decision-making with machine-learning, Agrawal, Gans and Goldfarb (2019), point to the difference between prediction and judgement, that is between calculating the expected payoff and likelihood of different outcomes (prediction) and determining the value of those payoffs in the first place, i.e., providing input to the decision-function (judgement). They argue that performance increases in AI are performance increases in prediction, leaving judgement as a human task.

Removing the steps of data analysis and prediction from human hands has clear advantages. Given proper training data, AI may remove some bias from corporate decision making that originates from the bounded rationality of human actors. At the same time, simply using AI far from guarantees the elimination of bias that is well-known to influence corporate decision making, such as hybris or politicking. Indeed, bias may emerge through strategic manipulation by the users themselves, as “agents strategically alter the input to the algorithm, perhaps because they stand to benefit from biased predictions” (Choudhury, Starr, & Agarwal, 2020, p. 1382). And on top of this, biases may be introduced unknowingly, such as when improperly trained models run the risk of becoming overfitted to poor training data, i.e., “producing results that are highly idiosyncratic to the observed sample” (Shrestha et al., 2020, p. 857). The call to managers is clear: Yes, AI can and likely will be an invaluable resource in data analysis for strategizing – but both machine- and human judgement-based inputs should be constantly questioned to ensure good results.

In the next steps of strategizing, the formulation of strategic actions and their eventual implementation, maintaining the right levels of human and AI responsibility will be equally important – although the balance might favor humans more in both cases. The translation between real-world impressions and data is a fundamental problem for complex AI applications. This issue, known as the frame problem in early AI developments (Shanahan, 2016), is problematic *once* for data analysis and forming recommendations; in codifying relevant factors of the real-world to allow an AI to work with the input. Once complete (to a satisfactory extent), however, both the actual process of AI calculation and the output generated by the AI stay in abstract the world of codified data. For strategy formulation and implementation, the translation is problematic *twice*; in translating real-world impressions into data before the use of AI algorithms, and again in translating algorithmic results back into implementable, real-world actions. This is where human intuition must once again play the dominant role – creating a system of augmentation that makes use of the strengths of both sets of actors (Haefner et al., 2021; Puranam, 2020).

AI and organizational structure

Key decisions on strategy should usually correspond to changes in organizational structure (Colfer & Baldwin, 2016; Donaldson, 2001). Re-shaping an organization from a product- to a service-oriented business model, for example, will bring with it a necessary restructuring of departments, processes, and responsibilities. As we have seen in the previous section, both strategizing towards- and with AI are extremely relevant to strategic leaders (see Table 3). For organizational structure, the immediate focus might lie on structuring *towards* AI –

determining systems that allow AI to be used effectively. Nonetheless, the act of structuring *with* AI tools deserves a second thought.

Software tools for the analysis of organizations are already in broad use and many related use cases may become more performant through the introduction of AI. Organizational network analysis, for example, in transitioning from analysis-based descriptive to AI-driven normative tools, could assist managers by providing outlines for improved department structures or required communication channels (Kearney, 2019). Allowing for organically floating, project-like structure, based on the recommendations of an AI tool, could allow for the easier distribution of information to relevant stakeholders in the organization. Considering the four essential tasks of organizations – task division, task allocation, provision of information, provision of rewards (Puranam, Alexy, & Reitzig, 2014) – such systems directly support, at least, the aspects of task allocation and the provision of information. In doing so, they begin automating, or at the very least augmenting, what previously was very much the domain of strategic leaders themselves. Use cases of this sort may become common, as access to AI tools becomes easier and organizations of all sizes may more easily begin experimenting. At the same time, the overall impact of these tools may be limited for the organization.

(Re-)defining a structure specifically *for* the efficient use of AI, however, will significantly impact the modus operandi of organizations. In their book “Competing in the Age of AI”, Iansiti and Lakhani (2020) outline what they call digital firms, as opposed to traditional firms. Digital firms, especially in their final stage, employing AI factories to deploy products as data use-cases, are defined by foundational data lake, to which all other areas of the organization have access. This allows rapid deployment of AI solutions in project-like structures. The authors see this as a departure from the “traditional, siloed structure of firms, which limits growth and responsiveness, prevents agile communication and coordination, localizes decision-making, and traps technology and data in isolated pockets”. (Iansiti & Lakhani, 2020, p. Preface).

Considering the four tasks of organizing described above, it is easy to see how strongly this transformation would impact the organization. The division of tasks, i.e., “mapping the goals of the organization into tasks and subtasks” (Puranam, Alexy, & Reitzig, 2014, p. 165), looks entirely different. Instead of being organized into departments with fixed, somewhat repetitive tasks, as is the case in most larger organizations today, employees now find themselves in use-case teams, working on an individual, integrated solution. Similarly, the allocation of tasks, i.e., “the problem of mapping the tasks obtained through task division to individual agents and groups of agents” (Puranam, Alexy, & Reitzig, 2014, p. 165), may look very different, as well. With new roles emerging through new skill requirements and the already existing roles re-shuffled into new team structures, most agents in an organization will find themselves with a changed, if not entirely new, set of tasks. An adjusted provision of information necessarily follows, which should ideally allow for much more efficient means of communication – remembering that one reason for the existence of the AI factory outlined by Iansiti and Lakhani (2020) is the attempt to establish an organizational structure better suited for the creation of data-driven products and services. The provision of rewards, both monetary and nonmonetary, is the one task of organizing that could remain least affected by the introduction of AI. While we already know of organizations with strict regimes of supervision, as described in the section on organizational efficiency above, where rewards might be more closely tied to measurable performance levels for employees, this is not a necessary consequence of the introduction of AI. Organizations could just as well choose to maintain their existing rewards structures while transitioning other aspects of the organization towards becoming an AI factory.

The analysis above describes the extreme case of working towards AI, where an organization decides to make AI part of their raison d’être. The natural follow-up question is:

Will all organizations transition to this stage of AI-integration? Using the most essential definition of an organization as “as (1) a multiagent system with (2) identifiable boundaries and (3) system-level goals (purpose) toward which (4) the constituent agent’s efforts are expected to make a contribution” (Puranam, Alexy, & Reitzig, 2014, p. 163), the answer is likely no – not every NGO, or small craft-enterprise will have to adjust their core business model accordingly. Considering medium- and large for profit organizations more specifically, we once again refer to Iansiti and Lakhani (2020, p. 3) who say they can “almost guarantee that no field of human endeavor will remain independent from artificial intelligence. In discipline after discipline and industry after industry, digital networks and AI are becoming pervasive, defining a new age for business and for all of us” – implying that at least some, if not all structural changes described above must follow, and requiring strategic leaders of these organizations to prepare accordingly.

AI and organizational processes

Having considered strategy and structure, the final piece of the puzzle of organizational form and conduct are the organizational processes and routines that bring to life the designs and plans created for and with AI. The category of process-design *with* AI entails mainly automation for organizational efficiency, already covered earlier. We will therefore first focus on the human side of processes *towards* AI; ensuring that all members of the organization are prepared and willing to engage with the new, algorithmic colleagues. The discussion here centers around algorithmic aversion and trust.

Algorithm aversion describes the phenomenon that people will trust a human’s expertise and forecast more than an algorithm’s or weigh human input more strongly than algorithmic input in a variety of scenarios (Dietvorst, Simmons, & Massey, 2015; Prahl & Van Swol, 2017). Researchers have linked the rise of algorithm aversion to lack of transparency in data and algorithms. The complexity of big data, with increasing velocity and variety of changing sources and data types (Janssen, van der Voort, & Wahyudi, 2017), makes understanding the inputs that influence a prediction difficult to understand already, which is further exacerbated by black-box algorithms leaving users in the dark as to how the results are produced (Lu et al., 2019).

On the other hand, Logg, Minson and Moore (2019) have shown through a series of experiments that laypeople may actually prefer algorithmic over human advice, in a phenomenon they call algorithmic appreciation. Interestingly, this algorithmic appreciation disappears in areas where the subjects consider themselves to be an expert and begin relying on their own input rather than a third parties' – to the detriment of their performance. From a psychological perspective, this is understandable; we believe in data, even if it contradicts experts, but believe in ourselves even more than in data, when we believe ourselves to be the expert. Another series of experiments has shown that the use of algorithms reduce more dramatically, compared to the use of human advice, after receiving a poor piece of advice – the penalty for mistakes is higher for algorithms than human experts (Prahl & Van Swol, 2017). To organizations, this is a problem – the people having to work with algorithms are exactly those experts that previously made decisions themselves and the level of use should continuously high. The question therefore becomes: How can we increase trust in algorithms, even amongst experts in the field?

The first possible answer is: Through explainable AI. With modern systems, especially multi-layered neural networks, algorithms are becoming increasingly powerful, but also complex and untransparent for users and stakeholders. Explainable AI targets the creation of systems where all inputs, outputs, and calculations are traceable and understandable to outside observers. Increases in the transparency of algorithmic systems often go hand in hand with decreases in performance (Hagras, 2018; Linardatos, Papastefanopoulos, & Kotsiantis, 2020). Nonetheless, the benefit of explainable AI seems to place it at the heart of a new movement of development, with large corporations such as Mercedes-Benz (2022) or IBM

(2022a) even putting the concept at the core of their development efforts for both internal and external use cases.

A second strong influence on usage patterns of algorithms is the extent to which users feel involved in the creation of the algorithms and therefore results. Giving employees the option to modify algorithms, if ever so slightly, already increases the extent to which algorithms are used. Dietvorst, Simmons and Massey (2018, p. 1155) describe this as "indicative of a desire for *some* control over the forecasting outcome, and not for a desire for *greater* control over the forecasting outcome" [emphasis in the original], tying in to the previously mentioned psychological effects of humans wanting their own perceived expertise to be valued. Kawaguchi (2020, p. 19) finds the same effect in his study of vending machine operators who have their assortments recommended by a management-imposed algorithm, stating that "integrating a worker's opinion makes the worker more likely to follow the algorithmic advice."

Other factors shown to reduce algorithm aversion are new domains where personal experiences are reduced, or vice versa, the perceived performance, expertise and experience of the algorithm are increased (Bigman & Gray, 2018). Interestingly, Ghasemaghaei, Ebrahimi and Hassanein (2018) show that higher data analysis competency of users increases the technology-augmented performance in decision-making based on large data sets. This could imply, combined with the previous findings, that domain experts, traditionally focused on personal experience, may recognize better performance in decision-making with AI tools after extensive data literacy training, which could then lead to an increased level of trust in data-driven tools.

In their literature review, Glikson and Woolley (2020) further differentiate between cognitive and emotional trust – the former being affected by the performance-related factors mentioned so far. Emotional trust, on the other hand, seems to be strongly influenced by the anthropomorphism of the representation. This is especially significant for robotic AI, where human likeness creates a higher level of immediate trust, but is also true for virtual AI: a visual representation of the AI, coupled with human names and other human-like features such as tangibility and immediacy behaviors positively influence the initial trust placed in the AI (Glikson & Woolley, 2020). The further trajectory of trust, however, is then more reliant on the level of performance.

A last point to consider in process design, in this case *with* AI, not *towards* AI, is organizational control, or the algorithmic supervision of employee performance. AI tools can be employed to promote efficiency gains in a system with increasing datafication of organizational oversight and control (Schafheitle et al., 2020). These systems of control are often more comprehensive, but also more opaque than previous forms of control, once again raising the issue of black-box algorithms that prevent a clear understanding of the resulting performance rankings – while also removing the personal connection between manager and employee that previously existed (Kellogg, Valentine, & Christin, 2020). Considering this dilemma purely from a performance-standpoint, AI systems are more detailed, more truthful, unbiased by personal preferences, therefore equally fair and detailed for every employee, which increases productivity and performance. At the same time, they may miss out on performance indicators that are more difficult to codify, such as interpersonal value for a team, and lead to reduced engagement and motivation – thereby reducing productivity in turn. Tong et al. (2021) show that these two performance effects co-exist and that being placed under algorithmic supervision is especially problematic for employees with shorter tenures in organizations. The authors hypothesize that these employees lack the emotional support system within the organization to counterbalance the negative performance impact of AI feedback.

All findings outlined above point to the importance of strategic (human) leaders in an organization employing AI in their internal processes. Finding the strategic middle ground in choices such as explainability versus performance, employee well-being versus automation, or individual control versus cost-efficient standardization may strategically be just as impactful as the decision to employ AI in the first place. Employees may frequently have clear ideas as to how their tasks and routines should be structured – with purely performance-oriented data perhaps pointing to a different route. Making employees feel respected and ready to engage with AI, while still reaping the expected organizational performance gains, might be the key challenge for strategic leaders in the coming years.

AI and the role, tasks, and tools of strategic leaders

In the final part of this analysis, we now fully focus on strategic leaders themselves – how do their roles and tasks change in the world of AI? Maintaining our distinction between organizing *towards* AI and *with* AI, we note how these become increasingly interconnected when looking at questions of strategic leadership. On one hand, leaders who only intend to use AI as a tool may nevertheless need to prepare their organization adequately to be able to select and deploy these tools effectively. On the other hand, even managers designing full AI-first organizations will face similar challenges in managing AI-related day-to-day activities as leaders of firms which only draw on AI as a tool. Accordingly, in the following we reflect on the role of strategic leaders following a chronological approach, beginning with the introduction of AI to the organization and then looking at execution, without implying that all firms need to go through the processes we describe in that order.

Table 4: Effects of AI on the tasks and roles of strategic leaders

Aspect of organizational form/conduct	Goal of AI usage **Towards AI** Introducing AI as organizational goal	 **With AI** Introducing AI as organizational tool
Tasks	Ensure adaptive capacity in organization – resource availability, flexibility, ability to adjust to new conditions, innovative culture	Fulfill requirements for ethical business decisions – ensure transparency and ongoing monitoring of systems
Roles / Requirements	Foster own absorptive capacity – willingness to learn and show digital affinity, adjust mindset, engage with opportunities	Allow augmentation of own adaptive capacity – AI support in personal areas of expertise, decisions

Strategic leadership towards AI: Introducing

In their review of strategic leadership research, Boal and Hooijberg (2001) point to two key tasks of leaders: Creating and maintaining both absorptive and adaptive capacity – absorptive capacity referring "to the ability to learn. It involves the capacity to recognize new information, assimilate it, and apply it toward new ends… [which] occurs at both the individual and organizational levels", and adaptive capacity referring "to the ability to change… [the] strategic flexibility [that] allows a firm to proact or respond quickly to changing competitive conditions" (Boal & Hooijberg, 2001, p. 517). Both the absorptive and adaptive capacity of organizations as well as their leaders will be tested throughout the entire process of recognizing and implementing AI.

We begin with the absorptive capacity of leaders – the ability to recognize new information and create a new reality, mental at first, in which AI becomes an essential part of the organization. This is the first required characteristic of strategic leaders: The capability of foresight, of imagining possibly radical changes to the status quo which might require overcoming significant organizational inertia. In the context of AI and digitization this might also entail a personal interest and willingness to learn about these technologies. In any organization of a significant size, it is likely that engineers or developers will be the first to play around with AI tools, out of educational expertise or personal interest. Leaders must then be willing to engage with these employees, display a certain level of digital affinity and openness to novelty in engaging with digital tools for optimal firm outcomes (Heavey et al., 2020; Merendino et al., 2018).

Being able to sense and understand the opportunities offered by AI therefore may rely on strategic leaders' willingness to adapt continuously. Being able to implement these opportunities effectively, however, is more a question of the level of adaptive capacity the leaders have been previously able to instill into the organization. Strategic leaders' success at such long-term efforts have been shown to correlate with their individual characteristics.

Research on managers' personalities and the links to performance in human-machine-interaction systems reaches back many decades, with Wynne and Dickson (1975) already linking success in commodity trading to personality traits such as defensiveness or need-achievement. Similarly, in the area of management, upper echelon theory highlights how organizations often reflect leaders' characteristics (Finkelstein & Hambrick, 1996), including the adaptability of a company (Liu, Fisher, & Chen, 2018) and the innovative culture of firms (Zhang et al., 2017). Similarly, biases arising from leaders' demographics, educational backgrounds, and track records, and from how leadership teams are composed are connected to technological innovation (Ahuja, Lampert, & Tandon, 2008; Carpenter, 2002; Finkelstein & Hambrick, 1996; Hayward, Rindova, & Pollock, 2004)

These studies all point to a similar finding: "Organizational flexibility derives from the leaders at the top." (Boal & Hooijberg, 2001, p. 517), requiring leaders to display not only an acceptance of change (Black & Boal, 1996), but also an active interest in preparing their organizations for change even when it is not yet clear when or how it may occur. This is the situation many strategic leaders find themselves in at the time of writing. It is clear that AI will impact the world of business – and even leaders in industries where products and services will only change marginally are recognizing the potential disruption coming their way in terms of organizing. Our call to action is therefore not one of forcing immediate change – instead, it is to ensure a level of preparedness within the organization. Resources can already be made available wherever possible, data can be cleaned and made available in the right formats, training programs can be created, etc. All this increases the adaptive capacity of the organization ahead of time, making the transition easier once the right AI use cases are found.

Strategic leadership with AI: Executing

In the active stage of an organization's AI journey, working with AI as an organizational tool, the requirements and expectations for strategic leaders change. We previously pointed to the absorptive capacity of leaders as a key requirement to recognize new opportunities.

Interestingly, augmenting this absorptive capacity is exactly where the strongest benefit of AI lies, on both the individual and organizational level. The capacity to handle large amounts of information, assimilate it into existing knowledge structures and apply it towards the most productive ends for the organization is what complex AI systems allow an organization to do better, faster, and more efficiently than before. On the level of the organization, this might take the form of strategizing with AI, outlined earlier, where tasks such as the analysis of complex market environments can be automated entirely. On the level of the individual leader, this means that some of the key decision made are no longer based on their managerial expertise, but may be augmented, for example in sequential AI-to-human decision making models (Shrestha, Ben-Menahem, & von Krogh, 2019).

From their review of academic literature in the field, Samimi et al. (2022) identify eight specific tasks of strategic leaders, allowing us to consider these implications in more detail. The tasks are making strategic decisions, engaging with external stakeholders, performing HR management activities, motivating and influencing organizational members, managing information, overseeing operations and administration, managing social and ethical issues, and finally managing conflicting demands.

Of these tasks, we have already seen how making strategic decisions, performing HR management activities, motivating and influencing organizational members, managing information, and overseeing operations can be heavily influenced by the introduction of AI. Engagement with external stakeholders could be added to this list through, e.g., algorithm-based personalized advertising to key customer groups, or public relation tools utilizing sentiment analysis to better understand the requirements of broader stakeholder groups. At the same time as AI gains prominence in these areas, however, we believe that the human influence of leadership is more important than ever. We have already described this with regard to the possible negative effects of relying on AI in motivating and influencing employees and overseeing operations and administration.

For the tasks of managing social and ethical issues and managing conflicting demands, the issue becomes even more pressing. Controlling unethical or illegal behaviors of a firm operating with extensive AI tools is only possible when the developers, users, and managers of the organization are aware of what and how the AI makes decisions. In supplier management, for example, an AI can be used to monitor a vast range of performance indicators of global suppliers – making thousands of predictions on price development, risk assessment, or supplier reliability every day (Dash et al., 2019; Toorajipour et al., 2021). However, the AI can only use what it knows, and so if certain, for example humanitarian, factors are excluded from the automated analysis, the AI could easily end up supporting suppliers with corrupt or dangerous practices. Similarly, Abada and Lambin (2023) show that, in scenarios where a limited number of agents all use simple and independent machine-learning algorithms to buy and sell a storable good, the machines appear to reach collusive price alignments. Such scenarios require intervention by regulators or managers to achieve more socially desirable outcomes. Explainable AI is important, but only part of the solution – the control of technology must reach deeper and include ongoing monitoring and critical questioning of the tools employed.

Managing conflicting demands is an area that will still require strong managerial involvement in the future. Agrawal, Gans and Goldfarb (2019) introduce a distinction between prediction and judgement in the age of AI, showing that, as AI becomes increasingly performant in predicting outcomes based on certain input variables and values, the judgement, i.e., the act of valuing those outcomes in the first place, becomes increasingly important for organizations. Only in this complementarity, finding the right type of conjoined agency with technologies, can organizations evolve to benefit from algorithmic systems (Murray, Rhymer, & Sirmon, 2021). Smith and Beretta (2021) name this struggle between autonomy and control as one of the emerging paradoxes of organizational structures; allowing transformative technological tools and the associated teams to operate to their highest levels of performance, while also maintaining a level of control that allows for transparency and efficiency across tools and teams – especially when competing demands arise.

In closing, we want to reiterate a point made in the section on the emergence of AI. The advantages of algorithms apply only to decisions made within the existing frame of an organization's operations. All algorithmic systems rely on training data to operate productively. This means that truly revolutionary changes, for example to an organizations' business model, will still require human input. Metaphorically speaking, an AI could not have predicted the rise of AI – and so, while individual tasks may be augmented or automated and therefore change the requirements of strategic leaders on a tactical, day-to-day basis, they still face the challenge of recognizing the grand strategic opportunities and risks before data is available.

DISCUSSION

In this chapter, we have laid out how the introduction of AI might impact all facets of strategic leadership; beginning with direct performance effects of AI on organizational efficiency and effectiveness, to potential adjustments in organizational form and conduct, to potential changes in strategic leaders' tasks and actions. A key theme emerging from our analysis is the question of using AI solely as an additional tool in the repertoire of organizational processes versus adapting core aspects of the organizational identity to revolve around AI. As we elaborate below, current research and practice mainly revolve around clear applications of AI such as process automation or recommender systems. We expect that strategic leadership should play a particular role as organizations try to build actual AI-first organizations in which firms may not only attempt to produce radically new products, but possibly also radically overhaul the organizations that produce them.

Managerial implications

For practicing managers, we present, beyond the practitioner guidelines we referenced, a series of directly implementable recommendations, which we summarize in Table 5. Most of these directly follow from our above analysis, with the exception of the upper left field. Here, we suggest that organizations hoping to become AI-first will need to ensure that their strategic leaders become sufficiently familiar with AI to be able to see more in it than just another automation technology. To escape this trap, "managers will have to retool and learn both the foundational knowledge behind AI and the ways that technology can be effectively deployed" (Iansiti & Lakhani, 2020, pp. 218-219). In this vein, several firms have rolled out broad trainings for their upper and middle management, the previously mentioned Microsoft amongst them.

Table 5: Overview of managerial implications of AI: Applications of AI and requirements for strategic leaders

	AI approaches	
Leadership aspects	**Towards AI** Introducing AI as organizational goal	**With AI** Introducing AI as organizational tool
Organizational performance		
Effectiveness / product innovation	Familiarize with AI to holistically understand use cases that go beyond replacing existing activities – and their impact on the workforce	Support incremental innovation in existing business models through improved data analysis in known search spaces
Efficiency / process innovation		Automate repetitive, time-consuming operations, primarily as cost-saving measure
Organizational form & conduct		
Strategy	Create new business models to make use of AI capabilities based on existing competencies, e.g., services, product features	Analyze markets, competitors, and consumers – predict most rewarding strategic actions
Structure	Adjust operations to suit data-driven outputs, e.g., project-based organizational structure	Analyze internal networks, e.g., for formalization of informal communication channels
Processes	Ensure acceptance of AI, based on trust in technology, likely employing explainable systems	Supervise operations and provide performance feedback – while ensuring acceptance and trust through balanced measures
Strategic leadership		
Tasks	Ensure adaptive capacity in organization – resource availability, flexibility, ability to adjust to new conditions, innovative culture	Fulfill requirements for ethical business decisions – ensure transparency and ongoing monitoring of systems
Roles / requirements	Foster own absorptive capacity – willingness to learn and show digital affinity, adjust mindset, engage with opportunities	Allow augmentation of own tasks, adaptive capacity – AI support in personal areas of expertise, decisions

Future research agenda

Like the technology itself, we expect the field of organizational research on AI to be highly dynamic over the coming years. Our analysis already suggests a series of questions scholars of strategic leadership may find interesting to inquire about to understand better how AI may impact organizational performance, form, and the role of strategic leaders therein.

Much of the growing collection of literature in the field, partly summarized in this chapter, discusses the processes and requirements of organizations beginning to work with AI. For the future, shown in Table 6, possible research questions will need to center around the effects of working with AI in the long run. In the area of organizational performance, this must include addressing how AI is embedded into larger technological and societal developments that continue shaping the world surrounding organizations, which may in turn define what AI is seen as and what it may be used for (Bailey et al., 2022).

Table 6: Selected directions for future research

	Future research areas	
	Towards AI	**With AI**
Leadership aspects	Introducing AI as organizational goal	Introducing AI as organizational tool
Organizational performance	How is AI embeddedness into larger societal and technological developments relevant to organizations	Can AI achieve, push, or support in radical innovation? What does the merge of local and distant search imply for organization theory?
Organizational form and conduct	How can leaders ensure long-term trust in augmented organizations? Which new forms of organizing can emerge through AI-based coordination? E.g.; -- How might AI help in scaling organizations with flat hierarchies? -- What are advantages and limitations of Decentralized Autonomous Organizations (DAO)?	How can AI strengthen organizational ambidexterity, e.g., to dampen the effect of external shocks? What are organizational and ethical limitations of algorithmic supervision and performance evaluation?
Strategic leadership	How can leaders ensure AI-specific adaptive capacity in the organization? What roles will leaders have in new forms of organizing, e.g., reduced emphasis on expertise vs. ethical guidelines, motivation, people-management? How can we ensure ethical use of AI across organizations, countries, continents?	Which characteristics of leaders make them likely to successfully adopt AI tools for themselves and the organization, i.e., ensure adaptive capacity – are there connections to, e.g., conscientiousness, control, strategic thinking?

Looking at organizational form and conduct, on top of further developing ongoing research around long-term trust in AI with stronger empirical results, researchers may begin looking at the emergence, scalability, and limitations of adjusted- (e.g. flatter hierarchies), or entirely new forms of organizing (e.g., DAOs) and how they impact our traditional understanding of leadership (Hsieh & Vergne, 2022). Similarly, shaping existing organizational and personal capabilities will become essential for leaders, with the adaptive and absorptive capacities likely determining the success of AI implementation. Researchers may therefore focus on determining how exactly these capacities may be shaped and strengthened.

A large space for possible AI-related research opens up around organizational innovation, for both theory and practice. With AI already taking over a stronger role in product development, we see this trend continuing over the coming years, posing questions of how AI might be useful in predicting or dampening the effects of external shocks in organizations. Researchers may therefore investigate both conceptually and empirically the processes of using AI to strengthen ambidexterity or radical innovation. In traditional organizational theory, some fundamental concepts may be challenged and require adjustments or clarification, e.g., what might a disappearance of local and distant search imply for the competitive environment and organizations struggling to identify valuable niches.

Spanning across all areas of research are increasingly important questions around the ethical use of AI. These questions address AI as a tool in established organizational applications such as the supervision of employees, but also in newly emerging scenarios, as leaders must face new responsibilities in an increasingly digital world where issues around privacy, data security, or social inequality will rightfully continue to gain prominence.

With every new generation of AI tools, likely more performant than the one before, practitioners and the academic community may need to reconsider the question of what these technologies enable and what the impact on organizations may be. To researchers of strategic leadership, on key question is whether an AI will ever be able to represent enough of the real-world environment of the organization to analyze and guide not only to incremental improvements, but to larger, perhaps even radical, shifts in business models and whether managers will still play a role in these decisions. At this point in time, we feel confident in saying that a future of AI entirely without strategic leaders is unlikely. Nonetheless, we encourage researchers to keep asking these questions, to keep monitoring technological improvements, and to keep pushing the boundaries of analysis in line with the ever-expanding boundaries of this exciting technology.

Chapter 2:
Organizing for AI – Multiple Goals and Structural Dynamics in the Introduction of a General Purpose Technology

---

Note:
Earlier versions of this chapter were presented at the 2023 Academy of Management Conference (Huber, 2023) and the 2024 Organization Science Winter Conference (Huber & Reetz, 2024). The version shown below is co-authored by David K. Reetz.

Organizational structure is central to how firms compete. Established structures, however, may be challenged by a changing technological landscape (e.g., Tushman & Anderson, 1986). A vibrant body of literature suggests that in such situations of rapid technological change, firms may benefit from designing constantly adapting rather than rigid organizational structures (Clement & Puranam, 2017; Joseph & Gaba, 2020; Rindova & Kotha, 2001; Siggelkow, 2002). Such dynamic structures promise to allow firms to profit from broad exploration to identify and test ways of implementing the new technologies, while also retaining sufficient coordination across interdependent activities at key moments to exploit synergies. Firms thus adapt by balancing decentralized exploration and centralized control (Siggelkow & Levinthal, 2003): while learning how to apply novel technologies in a value-accretive fashion they also adjust their structures to fit the emergent use (Burns & Stalker, 1961; Miller, 1996).

Considering the introduction of artificial intelligence (AI), as the latest significant change in the technological landscape of organizations, this adaptation of structures may be more challenging than before. On the surface, we see similarities to previous shifts; promising novel opportunities for organizations to improve efficiency and innovate, recent recommendation for practitioners focus on dynamic structures (Agrawal, Gans, & Goldfarb, 2020; Iansiti & Nadella, 2022) – mirroring the ideas presented in the literature above. However, we may consider AI to be not just a new technology, but the latest digital general purpose technology (Crafts, 2021; Goldfarb, Taska, & Teodoridis, 2023; Jovanovic & Rousseau, 2005), implying near-limitless possible applications, with major implications for organizations and entire industries (Agrawal, Gans, & Goldfarb, 2022b; Agrawal, Gans, & Goldfarb, 2024; Iansiti & Lakhani, 2020). This places organizations in a position of fundamental uncertainty; they must identify the strategy and structure required to achieve a certain outcome, but at the same will struggle to predict and evaluate these outcomes in the first place (e.g. Ehrig & Foss, 2022; Knight, 1921; Townsend et al., 2018). Put differently, AI may not only affect how organizations employ technology in their existing business and operational models, but also alter the underlying payoff structures (Gavetti, Helfat, & Marengo, 2017), as "profiting from enabling technologies [seems] fundamentally different from profiting from more narrowly applicable ones" (Gambardella et al., 2021, p. 76).

The fundamental uncertainty associated with AI suggests that decentralized exploration may be required to enable imagination of novel strategic options (e.g., Alvarez & Porac, 2020; Rindova & Martins, 2023) – the most promising of which often emerge as a result of individual actors' idiosyncratic representations (Felin & Zenger, 2017; Gavetti, 2012). With multiple actors pursuing exploration, organizations may find themselves confronted with multiple and conflicting goals (Gaba & Greve, 2019; Levinthal & Rerup, 2020) that stretch beyond their prior experience. But because this experience is imperative for the existence of shared mental representations that allow for the evaluation of feedback and implementation of change (Csaszar & Levinthal, 2016; Puranam, Alexy, & Reitzig, 2014)[6], lacking it will, in turn, impede coordination and alignment. It thus remains unclear how, in the face of fundamental uncertainty, firms manage to balance decentral exploration and centralized control.

[6] In line with this literature, we define shared representations as (implicit) perceptual mechanisms that allow interacting individuals to coordinate their actions. For example, to agree on of what AI 'is', what it may be used for, and how to measure the value of applications.

With this project, we aim to improve our understanding of how organizations manage the process of GPT introduction – balancing decentralization and control, and especially the shifts between them, to resolve fundamental uncertainty. We thus respond to a previously voiced shortage of research examining the implications of structure and information processing in the early stages of organizational decision-making (Joseph & Gaba, 2020). Building on a qualitative analysis of three multinational corporations, using archival data and primary interviews, we present a process model of GPT introduction. The model reveals distinct stages of organizing and highlights the importance of transition periods. We show how organizations manage to create shared mental representations from conflicting goals, through initial exploration followed by a standardization of processes and tight control (Boumgarden, Nickerson, & Zenger, 2012; Siggelkow & Levinthal, 2003). Rather than for exploitation, as previously theorized, we show how such standardization is essential to the creation of shared representations (Joseph & Gaba, 2020; Posen et al., 2018), which then allows organizations to use more decentralized, platform-like structures (Gregory et al., 2021; Kretschmer et al., 2022), to allow for the diffusion of technology and its use in specific domains. We discuss how such shifts in structure represent the dynamics of power and the development of understanding between each stage, as organizations move towards coordinated action.

LITERATURE AND MOTIVATION

All organizations face the ongoing struggle of identifying suitable activity configurations to successfully navigate their environment of stakeholders; employees, customers, suppliers, competitors, or regulators. This challenge is "particularly acute after environmental changes, such as technological shifts, that allow new ways of performing existing activities" (Siggelkow & Levinthal, 2003, p. 650). Technological shifts caused by the introduction of general purpose technologies (GPTs) may add to the significance of this challenge (Gambardella et al., 2021).

GPTs are defined as being pervasive, improving over time, and spawning further innovation (Bresnahan & Trajtenberg, 1995; Jovanovic & Rousseau, 2005). Technological improvements are most commonly narrow or clearly targeted at the improvement of existing activities, and therefore slot neatly into existing systems of configurations. GPTs, on the other hand, ask a more fundamental question of organizations as they effect broad systems change (Agrawal, Gans, & Goldfarb, 2024). The outcomes of their introduction into existing systems are, by definition, unpredictable ex ante, providing fundamental uncertainty to organizations.

With organizations uncertain of the possible uses of the technology at the outset, they must begin experimenting with solutions while not yet understanding the outcome they are looking for. This means that "the search for more or less appropriate strategies should incorporate exploring not only the space of alternative policies, but also the space of possible representations under which these policies might be evaluated" (Csaszar & Levinthal, 2016, p. 2032), that is, the payoff structures used to assess new configurations. Organizations are thus facing a dilemma of resource allocation as a key design choice, as they must balance investing in variation, that is the creation of new structures, and selection, that is establishing control and stability to evaluate these configurations. While creating variation allows organizations to expand the option space and thereby establish a broader range of opportunities to choose from, it is also costly. At the same time, stopping the process of variation prematurely to focus on evaluation and deployment can lead to inferior performance, as (Siggelkow & Rivkin, 2005, p. 102) observe: "in complex settings, firms need designs that permit them to search a diverse array of operational configurations before locking in on a set of choices".

In terms of organization design, this dilemma often emerges as a choice between the decentralization and centralization of structures and decision-making authority, where we use the term decentralized to describe an organization structure in which "decision making has been disaggregated into a number of subunits, or divisions […] In contrast, an organizational structure is called 'centralized' when decisions are made only at the level of the firm as a whole" (Siggelkow & Levinthal, 2003, p. 651). Decentralized structures allow departments or individuals to exploit idiosyncratic knowledge and focus on their own best interests. They may choose to research issues bothering them specifically, to focus on the KPI against which their individual performance is measured, to use their human and monetary resources in ways they see fit, or, in the language of AI, to develop tools and use cases most relevant to their specific domain.

Centralized structures, on the other hand, most commonly characterized by management boards and various headquartered teams, focus on the organization as a whole. They help ensure that individual departments are aligned in their plans and approaches and ensure that success stories are communicated effectively so that others may learn from them.

Existing literature shows a number of approaches as to how organizations may manage the emerging dilemma, which allows them to "one hand, search broadly for new activity configurations [and] on the other hand […] coordinate across their interdependent activities to avoid misfits and instability" (Siggelkow & Levinthal, 2003, p. 651). For complex systems of nondecomposable activities, defined as those that cannot be completely separated but instead require some cross-divisional interaction and which likely arise as new technologies are introduced, Siggelkow and Levinthal (2003, p. 652) find that "neither a centralized nor a permanently decentralized organizational structure leads to high performance. In this case, temporary decentralization with subsequent reintegration […] yields the highest long-term performance".

Other authors have used similarly dynamic concepts to describe how organizations may effectively maneuver technology transitions: Boumgarden, Nickerson and Zenger (2012) use the metaphor of sailing into the wind to describe idea of vacillating. In this metaphor, organizations recognize that there is no clearly superior strategy available to them, but that they must instead accept some inefficiencies and regular course corrections. While "each course correction or 'tack' imposes a loss in forward momentum, the skillful sailors masters these reconfigurations so as to minimize momentum loss” (Boumgarden, Nickerson, & Zenger, 2012, p. 606). Along a similar line of argumentation, Smith and Lewis (2011, p. 393) suggest paradoxical elements in organizing best be resolved by “purposeful iterations between alternatives in order to ensure simultaneous attention to them over time. Doing so involves consistent inconsistency as managers frequently and dynamically shift decisions. Actors therefore make choices in the short term while remaining acutely aware of accepting contradiction in the long term.”

Fang, Lee and Schilling (2010, p. 627) describe a similar organizational design with their concept of semi-isolated subgroups, where, by “decentralizing the learning process to subunits of the organization and providing barriers to the rapid diffusion of ideas and norms across those subunits, managers may be able to encourage the exploration of a more diverse range of solutions”. The remaining moderate amount of connectivity to the network still enables “exploitation by facilitating the rapid diffusion and assimilation of currently superior knowledge” (Fang, Lee, & Schilling, 2010, p. 636), once identified in the process of isolated learning. Knudsen and Srikanth (2014, p. 414) mirror this sentiment, suggesting that “a high level of information transfer is unnecessary and perhaps even harmful in coordinated exploration”.

Jointly, this literature provides tremendous insights into the options of organizations aiming to find balance between exploration for new opportunities and identifying the most rewarding use of technology in existing structures.7 This line of work may help firms in transitioning to many new technologies. Yet, we suggest that when the technological change in question is a GPT, particularly a digital one like AI, where organizations deal with the

[7] The usefulness of these recommendations is echoed when turning to literature focused on practice – for example, work on how to best introduce AI to organizations also frequently recommends staged approaches, in which organizations iterate between decentralized structures allowing for exploration of valuable use cases, and centralized structures to scale use cases throughout the organization (Iansiti & Nadella, 2022; Kruhse-Lehtonen & Hofmann, 2020; Leonardi, 2020). Notably, though, we suggest that even the practice-focused literature seems to make the same the same assumptions we challenge next.

fundamental uncertainty of unknown unknowns on possible future states, these findings may fall short. We thus shine a closer light at some of the key assumptions this body of work appears to make and discuss how these may no longer apply with regard to AI. Particularly, we challenge the three notions of (i) shared mental representations between all involved actors, (ii) attaining and instantaneously evaluating performance feedback, and (iii) the readiness and capacity to quickly transition to new states of organizing.

The assumption of organizations holding shared mental representations is questionable in the context of GPTs. Mental representations, made explicit by Knudsen and Srikanth (2014, p. 427), are "models of reality held in the mind of an individual, who can use this representation to generate predictions about reality", which allow managers and organization designers to "consider the merit of alternative strategies without the need to actually invest in and carry out the various options", which must also include "the space of possible representations under which these policies might be evaluated" (Csaszar & Levinthal, 2016, p. 2031, see also Craik, 1943). Such representations are the foundation of organizational structures. Building on a joint belief of, for example, the importance of product lines, or the flow of work processes, organizations shape their hierarchies and departments. Furthermore, mental models can have significant effects on individual behavior, determining where and how much attention is paid to certain ideas: As they are necessarily limited images of reality, "the more the mental maps are aligned, the more the agents influence each other in concentrating on a narrow portion of the landscape that both see as beneficial… at a cost [of] a more superficial understanding of other regions in the landscape that perhaps are more valuable" (Knudsen & Srikanth, 2014, p. 429).

Considering the introduction of GPTs, there may not be an alignment of mental representations. First, there may not even be clear representations to begin with – a clear understanding of the possible future realities and the results of various strategies is simply not available to actors dealing with new technologies that so fundamentally challenge established ways of working. One possibility of how organizations can reach alignment is through a dominant coalition (Cyert & March, 1963; Simon, 1957). Yet, in the phase of a newly emerging GPT that contains the opportunity to redesign (agentically and strategically) existing organizational routines and processes, the power of the existing dominant coalition may be contested by various actors (March, 1962; Nelson & Winter, 1982), if it even stretches to the new domain. Thus, it is not clear if and how, at the organizational level, a shared mental representation should emerge that would guide the firm in designing and adjusting the organizational structure guiding its introduction of AI.

This also affects the second notion presented above; the possibility of instantaneously evaluating performance feedback. Model-based work, for example, draws on this assumption to predict an organizations' developments over extended periods of time, assuming that "if the firm finds that the alternative has a higher performance value than the status quo […] the firm will implement the alternative" (Siggelkow & Levinthal, 2003, p. 654). But when firms have no experience with rich GPTs like AI, we see two problems with this choice. First, any organizational intervention will be temporally decoupled from eventual feedback. Identifying and building an organizational structure to support AI may take significant time – and observing and evaluating its performance afterwards will, as well. This delay is exacerbated by the fact that systemically embracing a GPT may often be a non-decomposable task (Siggelkow & Levinthal, 2003), with performance feedback on piece-meal adaptions only allowing, at best, for little inference toward which of a potential series of organizational configurations the organization should choose. But, second, even if feedback were to be received in real time, given the organization would likely be experimenting in parallel with various options to deploy novel strategies and structures, learning which options are better individually and, in particular, as complementary configurations, would remain a trying endeavor (Siggelkow & Rivkin, 2009).

These considerations directly link to the third notion: that firms may smoothly transition between various organizational structures they need to deploy over time, that is, that they may easily switch between decentralized exploration and centralized exploitation when they somehow identify that the other approach would promise superior performance. While the literature acknowledges that actors “face ambiguity about whether and how to update their own beliefs” (Fang, Lee, & Schilling, 2010, p. 631), eventually agents are assumed to be “profit seeking and therefore switch to the most promising alternative they currently perceive” (Knudsen & Srikanth, 2014, p. 421). Our first two challenges already imply that, at the organizational level, such a simple switch should be unlikely when mental models differ so that various actors would arrive at different interpretations of what the most promising alternative should be – if the surrounding ambiguity allows them to arrive at such an evaluation. But as alluded to already in our last point, even if the organization were able to arrive at a somewhat objective evaluation of what a good next step would be, implementing it would not just consume money (which may be priced in) but simply requires time. During that time, however, fundamental uncertainty means that new information may arise at any point that could change the original evaluation. As a result, we expect that decision makers should struggle significantly in identifying when, how, and why to commit to a different organizational structure.

In sum, there is no reason to believe that even with GPTs, an organization trying to introduce new technology should not iterate between centralization and decentralization. Yet, just how the organization would achieve this is unclear; our common understanding of the process of creating shared representations and enforcing organizational change, especially the transitions between types of structure, does not apply in situations of fundamental uncertainty. Accordingly, drawing on the example of AI, in this paper we ask: How exactly do organizations create the shared mental representations that then allow them to design and adapt the organizational structures required to introduce a groundbreaking GPT?

DATA AND METHODS

Research Setting and Design

To generate new theoretical insights and answer our research question (Edmondson & McManus, 2007) we conducted an inductive multi-case study(Eisenhardt & Graebner, 2007). This design allows for a comparative analysis and replication logic of emerging insights (Eisenhardt, 1989). Addressing our research question required in-depth access into an emergent phenomenon (Langley, 1999), that is, an organization's past, ongoing, and planned future strategies as well as organizational processes related to introducing AI.

As part of a larger research program investigating the emergent use of AI tools in organizations, we originally worked with 15 different organizations, for a total of 77 interviews, workshops, and discussion rounds over 2 years. Through these conversations, we obtained a deep and thorough understanding of how organizations of various sizes, structures, and industry backgrounds begin their AI journeys. From this large dataset, we used purposeful sampling (Glaser & Strauss, 1967), to identify six suitable organizations; all industry-leading multinational corporations to ensure comparability of structural backgrounds and competency and with sufficient history and breadth of AI use in the organization. We focus on larger organizations here as they provide sufficiently complex structures and promise a larger diversity of goals for AI application, suggesting that coordination, as our primary concept of interest, could be a salient issue. Further, their roles as industry leaders increased the importance of mastering the implementation of a promising technology such as AI to maintain their competitive position.

Our research question requires access to highly sensitive data, concerning ongoing decisions about the future use of a technology still relatively new to the firms under investigation (Brynjolfsson, Rock, & Syverson, 2019; McElheran et al., 2024), revealing firms' active sensemaking with C-level relevance and firm-wide impact. Therefore – while remaining in contact with the remaining three organizations from the larger sample, using them for constant comparison in the development of the model and as sparrings partners in workshops and discussion – we then chose to further zoom in on the three specific organizations. These three, which we refer to as Blue, Green, and Red, are the organizations that allowed us the best field-level access across hierarchy levels and departments, had similar timelines of AI use with their journeys beginning roughly 10 years before the start of data collection, while also promising to provide us with particularly in-depth insights from their variation in approaches and goals of AI use. Accordingly, we refer to informants as B#, G#, R#; with numbers indicating the order of interviews. The three organizations not selected for the final sample are anonymized as Purple, Yellow, and Green. Table 7 provides an overview of all available data sources for the six organizations, with additional details on the roles of interview informants.

Data Sources

For each of our focal case organizations, we draw on different data sources. We began by collecting publicly available data, identified using common search engines as well as the online archives of business magazines. Relying on key words such as "[company name] + artificial intelligence", "[company name] + digitization", or "[product name] + strategy", we were able to obtain information reaching back to the very beginning of the companies' AI efforts. These included press articles and media interviews given by executives, as well as publicly accessible firm documents like annual reports and press releases. Second, we accessed online networking portals such as LinkedIn, to check personal profiles of key personnel, confirming dates of entering and (where applicable) exiting the companies. Third this, we were able to access company-internal archives, including proprietary and confidential strategy documents and reports – which we were granted access to after building rapport during interviews with firm representatives. Combining these external and internal archives, we created archives of data equaling about 80 pages of detailed and partly exclusive information per case.

Table 7: Overview of data sources and informant background

	Case	Available Data Sources	Interview informant role	Department	Joined [1)]
Primary cases	Blue	Annual reports Press releases Media reports Interviews	B1: Executive Manager	Global data strategy	y_0
			B2: Product Owner	AI center of excellence	y_4
			B3: Senior Manager	IT strategy & support	y_{0-12}
			B4: Expert	Division AI strategy	y_2
			B5: Executive Manager	Division AI strategy	y_3
			B6: Manager	Service – Data strategy	y_{0-11}
	Green	Annual reports Press releases Media reports Internal documents Interviews Workshop participation	G1: Global Lead	AI strategy & development	y_1
			G2: Technology Lead	AI center of excellence	y_{0-6}
			G3: Program Manager	Global supply chain	y_{0-9}
			G4: Team Lead	Division supply chain	y_{0-14}
			G5: Head of Innovation	Business function	y_{0-1}
			G6: Team Lead	Business function	y_{0-15}
			G7: Technology & Project Lead	Business department	y_{0-13}
			G8: Process Automation Manager	Business department	y_{0-1}
	Red	Annual reports Press releases Media reports Interviews Presentation at firm roundtable Workshop participation	R1: Vice President	Platform data analytics	y_0
			R2: Director data science	Platform data analytics	y_4
			R3: Team lead analytics	Platform data analytics	y_0
			R4: Vice President	Core data analytics	y_5
Corroboration cases	Purple	Annual reports Press releases Media reports Interviews Workshop participation	Board member	Digital Strategy	y_{0-16}
			Vice President	Digital Strategy	y_{0-14}
			Vice President Sales	Business unit	y_{0-25}
			Vice President HR	Business unit	y_{0-11}
			Director Sales	Business unit	y_{0-31}
			Digital Product Owner	Business unit	y_{0-16}
			Director Transformation	Business unit	y_{0-9}
			VP Additive Manufacturing	Business unit	y_{0-9}
	Yellow	Annual reports Press releases Media reports Workshop participation	Global Head	AI Strategy	y_{0-2}
	Orange	Annual reports Press releases Media reports Interviews Workshop participation	Director Sales Excellence	Regional business division	y_{0-10}
			Director & Project Manager	Regional business function	y_{0-14}
			Head	Regional business division	y_{0-4}
			Account executive	Regional business division	y_4
			Sector executive	Regional business division	y_2

1) Year of joining case company. Negative years indicate joining company before the start of focused AI initiatives

To deepen our understanding of the organizational dynamics underlying these journeys and add deeper insights to decisions made at the time, we complemented these archival sources with interviews. Despite having established initial contact with all organizations through our personal networks, identifying and getting access to suitable informants, willing to provide the necessary depth of information, required additional efforts. These informants were high-level executives and operational employees who were deeply involved in the previous and/or ongoing decision-making related to the implementation of AI. Relying on a system of internal referrals, we asked our initial contacts to introduce us to other suitable candidates, which was particularly useful in bridging structural gaps to other relevant departments and colleagues. This snowballing approach (Biernacki & Waldorf, 1981) allowed us to ensure that we included different levels of hierarchy and expertise, ranging from corporate heads of technology to users of AI tools, and representing the distinct perspectives of corporate headquarters as well as business teams. We continued in this snowballing fashion until we were confident that we had identified informants from all relevant areas, who were highly knowledgeable about and actively involved in introducing AI in their companies (Suddaby, 2006). As such, we conducted 24 in-depth interviews (Spradley, 1979) with multiple informants from each of the case organizations, lasting between 40 and 90 minutes.

All interviews were conducted by the first author as intentionally open-ended conversations, inviting respondents to articulate events in their own words and make sense of what they perceived to be important steps, as well as their current related activities and emerging ideas (Spradley, 1979). The interviewer thus aimed to understand the informants' points of view before asking clarifying questions. Specifically, these aimed at triangulating each organizations' AI journeys as they emerged from insights provided by different informants, and revealed differences between departments and individuals, as well as with respect to our archival data. All interviews lasted between 40 and 90 minutes and (amid Covid-measures) were conducted virtually via the videoconference service Zoom. For all interviews, we created extensive hand-written notes, including direct quotes, that were typed up and structured into detailed records within 24 hours of the occurrence of an interview to mitigate recall bias (Ozcan & Eisenhardt, 2009). We were also allowed to fully record 16 of these interviews, which we transcribed verbatim and returned to our respondents for explicit confirmation of direct quotes.

Data Analysis

Consistent with common practice in inductive research, we combined analytic steps that were most suitable to address our research question (Pratt, Sonenshein, & Feldman, 2020).

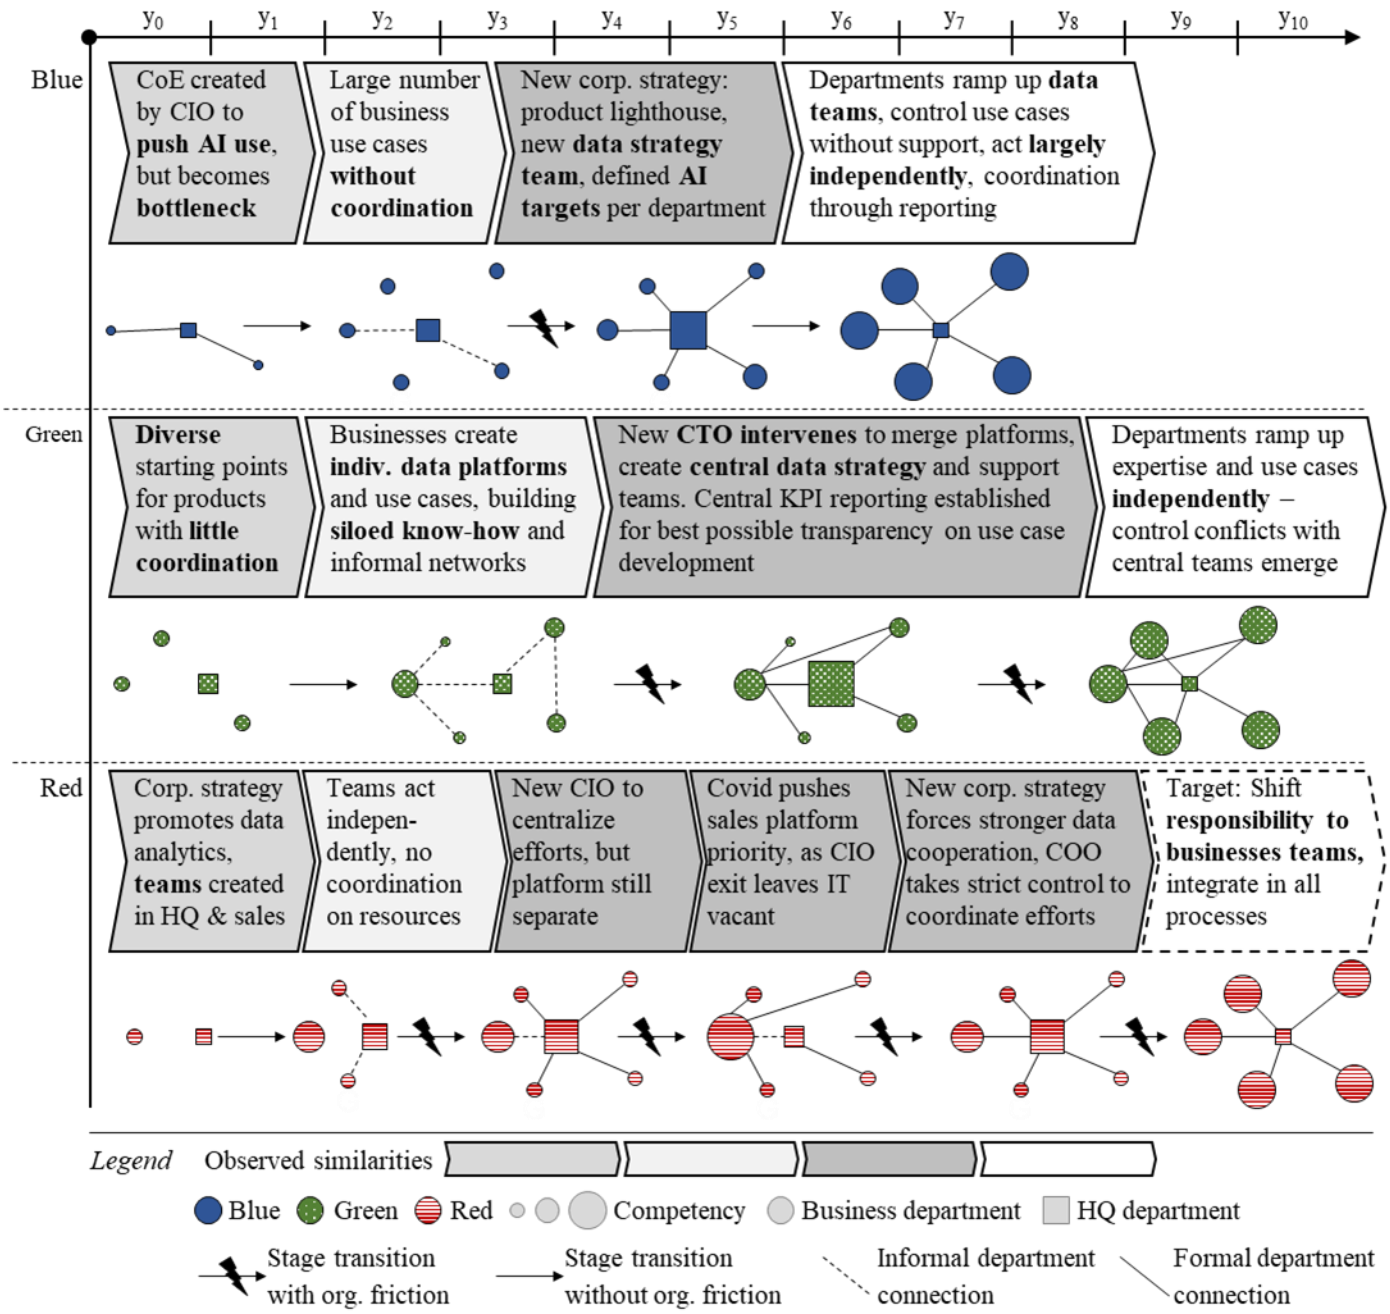

Figure 3: Timeline of cases with key events and organizational network diagrams

Constructing chronology

We began by using our archive of public data for each case organization to create initial chronologies of key events related to Blue, Green, and Red's introduction of AI (e.g., Gao &

McDonald, 2022). We then compared these insights with adjacent information related to the firms' broader product development, leadership changes, and engagement with their market environment, and treatment of previous technological changes – which allowed us to relate the implementation of AI to other firm activities as well as to delineate its specifics and the commonalities shared with previous events.

The resulting narratives of 13-15 single spaced pages per firm focused on the design or adaptation of organizational structure in general – marked by strategic announcements decisions – those committing resources or setting organizational precedents with large movements of resources (Mintzberg, Raisinghani, & Théorêt, 1976; Ott & Eisenhardt, 2020). These case histories provided us with a general understanding of each organizations' respective AI initiatives and strategies and served as an important basis for our further analysis (see Figure 3 for an overview).

Exploring concepts

Drawing on all our data sources for each case, we openly coded (Strauss & Corbin, 1990) how the firms set up organizational structures – comprising aspects such as ongoing projects, internal coordination, management support, distribution of responsibilities, collaborations, political dynamics, communication across departments, formal and informal networks, future initiatives, and visions of how to work with AI. We aggregated the resulting concepts of organizational activities regarding AI implementation by identifying common themes. Tables 14-18 in the Appendix provide the detailed evidence of this analytic step for each case.

By comparing these themes across our three focal cases (Eisenhardt, 1989), we noticed a consistent pattern: distinct altering stages of centralization and decentralization and significant transitional phases between them. We treated each of the stages and transitions as temporal brackets – that is, "comparative units of analysis for the exploration and replication of theoretical ideas" (Langley, 1999, p. 703). Despite their idiosyncratic starting conditions, all firms first formed decentralized structures, shifted to centralization, before moving back to decentralized ones, connected by phases of transitions.

Table 8: Data structure and emerging Temporal Brackets

Concepts	Themes	Temporal Brackets
Product-side reaction required to external developments	Observing external industry dynamics	Diverse Initiation
Recognizing the value of data through consumer-industry hype		
Leadership recognizes opportunity	Initiating internal technology push	
Create support structures with top-down funding		
Resource constraints limit reach	Arising organizational limitations	*Transition 1*
Urgency necessitates independent action		
Lacking buy-in through central push		
Available budgets for experimentation	Growing decentral curiosity and opportunity	
Bottom-up instead of corporate program		
Parallel structures without coordination		
Wild west of data exploration	Searching broadly for applications	Exploration
Departmental profile-sharpening		
Individual area-specific lighthouses		
Department-specific use cases, tech stacks, and processes	Driving local resource ramp-up	
Duplication of efforts with individual solutions		
Complex set-up emerging organically		
Waste of financial and human resources	Growing unrest with lack of outcomes	*Transition 2*
Growing disillusionment with lacking return on invest		
Need to separate the good from the bad		
Strategic decision towards data-driven organization	Increasing strategic attention	
Committing to AI as future focus technology		
Silos hinder cooperation – competition for en-vogue topics	Reshuffling political responsibilities	
Changing C-level appointments shuffle key departments		
C-level vacancies disintegrate entire teams		
Specify top-down AI targets	Agreeing on goals	Standardization
Define long-term crossfunctional strategy		
Collect toolbox of best-practice examples	Creating supporting structures	
Create end-to-end expertise in structures to grow AI literacy		
Provide central implementation support and regular exchange		
Prioritize high-impact projects with top-level mentorship	Establishing clear reporting lines	
Enforce KPI reporting and conversations across hierarchy		
Actively limit decentralized ramp-up of expertise		
Cooperation with central units limits speed of development	Emerging organizational limitations	*Transition 3*
Businesses now push IT for support, not the other way around		
Focus on domain knowledge, not central development		
Establish software-focus as long-term goal	Accepting loss of control	
Accept loss of control; trust is the new control		
Decentral competency to question central raison-d'etre		
Propose own ideas instead of waiting for IT	Expanding domain expertise	Diffusion
Continue building domain expertise to reduce risk of failure		
Domain-specific value creation as essential foundation		
Establish AI as standard practice	Formally integrating AI	
Machine learning has reached maturity, stop pushing		
Data analytics as business as usual in every division		
Become more vertical, with functional teams	Moving towards tech-first structures	
Decentralization as sign of sustainable success, not failure		
Coupled responsibility for business and AI		

Accordingly, in another round of coding, we focused our attention on the antecedents and characteristics of these transitions, especially with respect to understanding potential limitations of a respective organizational structure in advancing AI implementation. Table 8 presents the emerging data structure resulting from these analytic steps, with Figure 4 showing observed process dynamics.

Developing a conceptual framework.

Drawing on this temporal bracketing, we finally integrated our data structure, connecting stages and transitions, into a process framework of AI introduction (Miles & Huberman, 1994). Cycling back to literature on organization design and technology adoption (e.g., Gambardella et al., 2021; Siggelkow & Levinthal, 2003), as well as AI-specific literature (e.g., Agrawal, Gans, & Goldfarb, 2022a; Brynjolfsson, Rock, & Syverson, 2019; Frank et al., 2019; Iansiti & Lakhani, 2020; Iansiti & Nadella, 2022) we refined the emerging logic, vocabulary, and reasoning.

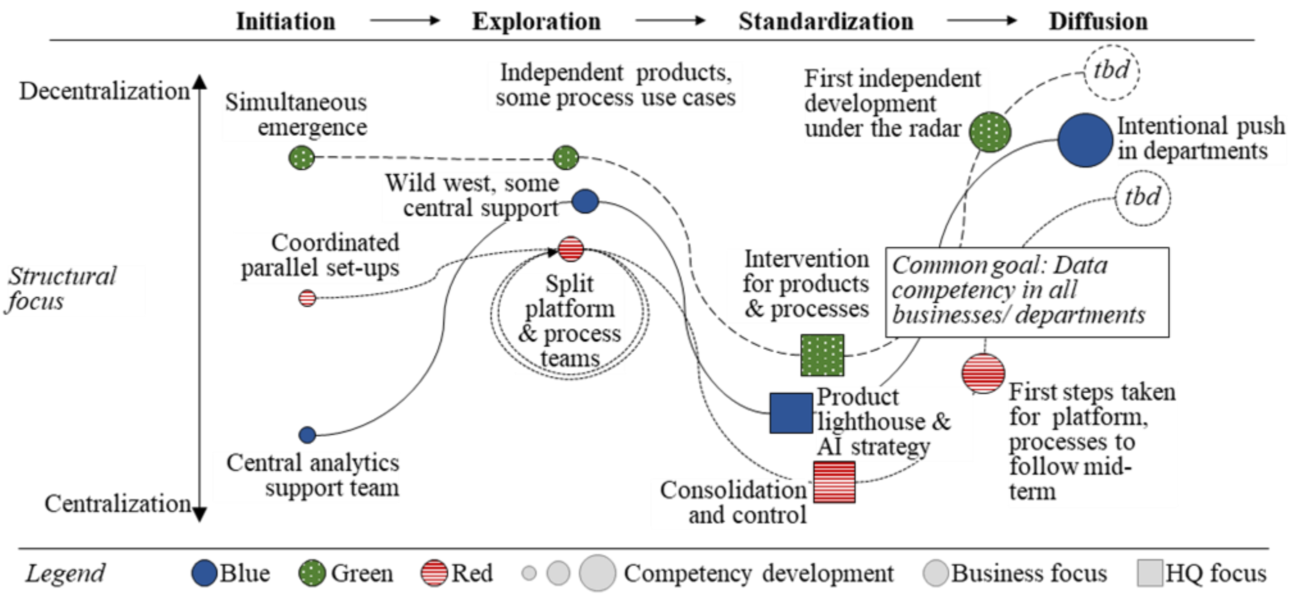

Figure 4: Case-specific dynamics of (de-)centralization across Temporal Brackets

We corroborated our emergent model in various ways. First, as noted above, we remained in contact with informants from other organizations in our initial sample. This allowed us to test key conceptual ideas, confirming that they are also reflective of their AI journeys. Second, we were invited to discuss our emergent findings with representatives of Red's and Green's data science communities. Through this iteration, we confirmed our timeline and organizational charts, as well as key concepts of the analysis. Finally, we organized a full-day workshop that was attended by three informants from our focal cases and four from our extended sample, as well as a consulting firm we had previously collaborated with around topics of AI. Presenting our preliminary findings and subsequently discussing them in detail helped us clarify and elaborate core themes of our analysis, as well as to strengthen the emerging parallels across cases.

Case Overview

The selected case companies, Blue, Green, and Red, are established organizations, global leaders in their respective fields, operate in mature markets, rely on established processes and structures, and possess significant experience in the use of business intelligence tools. The AI journeys of all organizations began at roughly the same time, between the years of 2012-2015, meaning that at the time of the interviews, all three were already several years into the process of transforming their operations towards a stronger focus on algorithmic data analytics.

In the following sections, we will briefly describe the dynamics of AI in each of the organizations as they may have been observable to the outside at any point in time – akin to our initial research on organizational timelines – before providing a more in-depth analysis of the dynamics in the following section.

Case background: Red

Red employs a B2C business model in one global market sector, outsourcing production to external manufacturers, and selling products through an integrated system of 3rd party retail, own stores, and an online shop. With production outsourced, the complexity of this business model is primarily driven by the anticipation of market developments in rapidly developing trends, as well as the complex logistical network of managing production and distribution. As basic over-the-counter goods, there is little AI relevance to Red's products. However, with a sales platform managed as key distribution channel and the logistics of sourcing products as a second use case, internal data analytics is gaining prominence. With little previous know-how, competencies are yet to be developed, although a straightforward business model and selected key applications that can drive significant organizational interest, could lend itself to a direct path of implementation.

Red's analytics operations were triggered by a central push consisting of three strong measures. First, communication in the y_0 annual report[8] promised significant improvements in the speed of internal decision making through data analytics. Second, the IT department was combined with supply chain under the board division operations, aiming to (as promised in the annual report) leverage real-time data analytics in everyday business processes. Third, a separate data team was created specifically for the B2C online sales platform.

Over the following years, Red underwent a series of leadership and structural changes. With both the COO and CIO leaving the company in short succession, the IT unit and its central data team was entirely moved into the CFO board division, also installing a new CIO. This new CIO then left the company again shortly afterwards, around the same time that the effects of the Covid pandemic hit – which forced the company to divert all efforts away from the central data team towards the sales platform analytics. Finally, in y_6 a new corporate strategy program was unveiled. The strategy announcement in the annual report made clear that the company wanted to strengthen the link between digitization initiatives and operations, looking to increase both speed and agility of the value chain. With these developments, the analytics initiatives were once again moved, this time to the new COO division, where they remain at the time of the interviews – striving to integrate analytics into everyday business processes throughout the organization.

[8] To ensure the case firms remain anonymous, we code years relatively to the year in which the firm introduced AI officially for the first time, which we define as year zero (y_0). Notably, across organizations, y_0 is different.

Case background: Blue

Blue also focuses on one primary industry category, integrating both B2C and B2B approaches in their hardware-driven business model. As an added layer of complexity, Blue controls its own material-intensive production, relying on an elaborate system of suppliers and manufacturing locations to serve global markets. This moderate level of business model complexity is paired with a moderate level of AI relevance for products, as the industry was expected to shift towards AI-based products in the medium-term, but without immediate key applications. Similarly, there was a medium level of AI relevance to internal processes, as aspects such as logistics or demand management promised to turn more data-driven in the future, but without any major effects on current processes expected. A series of highly successful business years provided sufficient organizational slack and resources to support a long-term investment, promising strong results.

While individual engineers in Blue's technical research departments already began experimenting with knowledge graphs and large-scale data analysis in the 1980s, the organization's AI journey began in earnest when a central data analytics unit was created in y_0. The foundation of the team coincided with and was aided by a favorable economic position, providing sufficient slack to prepare for potential future challenges. The new team was set up as an internal provider, hosting a local platform for data analytics services, as well as ideation, data strategy, and use case implementation support. Over the following years, however, individual business departments began implementing their own use cases without coordination, making little use of the central services and driving their own strategies instead.

Similar to Red, the introduction of a new corporate strategy provided a significant demarcation point towards the professionalization of Blue's analytics efforts: by y_3, the competitive landscape had shifted, and so Blue presented a new, AI-focused product strategy, thereby also pushing a more comprehensive approach to data analytics for the rest of the organization. The associated installation of a new central digital strategy team with the mandate to set AI targets and coordinate efforts across departments symbolizes this shift in approach. As departments began increasing their own data competencies over the following years, however, the focus of agency once again shifted – and by y_7, units appear to operate independently, in a system of distributed and democratized access to AI expertise, with usage starting to make its way into everyday processes across the organization.

Case background: Green

Green, the most complex of our three organizations, is divided into several business divisions, serving entirely different industries. The respective products include hardware- and software solutions for B2B customers, produced in own manufacturing locations. Similar to Blue, Green had organizational slack available at the beginning of their AI journey. As software products and long-term service contracts on their hardware products are major revenue drivers for Green, there is a direct AI relevance to their product offering – also building on strong existing software development competency in various division of the organization. With such strong interest in AI development, supported by available resources and previous expertise in the field, Green seemed to be the most promising of the three case companies at the outset of their AI journeys.

As such, it is no surprise that that AI engagement emerged organically across the organization, with most research going toward services such as data platforms for remote machine supervision. In fact, between y_o and y_4 multiple business divisions had already begun selling AI-driven service offerings to their customers, building on independent software developments. The approach changed in y_4, when a series of C-level decisions, aided by the implementation of a new CTO, introduced two clear measures for stronger centralization and coordination. Firstly, the most developed of the platforms currently being developed was selected to serve as the unified base-technology for all further applications, both regarding internal processes as well as external product offerings, extended through a range of APIs for other units to plug their solutions onto the existing stack. Secondly, the new CTO selected AI as a key future technology, providing funding for both a central data strategy team and an AI-specific corporate incubator. This led to a period of growing engagement with AI across the organization, marked by both informal structure such as exchange forums and educational formats, as well as formal structures such as clear reporting lines.

Around the time of data collection, Green seems to be finding itself at yet another interesting organizational crossroads. On one hand, internal and external communication strongly emphasizes the role of central functions, pushing an aligned strategy and the resourcefulness of the central teams. On the other hand, businesses seem to be ramping up their own, independently operating data science teams, and begin working on use cases without involvement of central support – thereby also mirroring the trends observable in the other case organizations.

FINDINGS: A CONCEPTUAL FRAMEWORK OF GPT INTRODUCTION

The conceptual framework emerging from this analysis reveals how the three focus case organizations, starting from highly idiosyncratic conditions, eventually evolved along similar paths to create shared representations from originally diverging interests. We show how actors attempt to strategically direct the collective attention of the organization, pursuing their individual and sometimes conflicting goals, and how this results in organizations where agency in the form of decision-making authority swings between centralized and decentralized structures. We further show how all organizations end up using seemingly stable, platform-like structures and how these can once again be challenged by further developments in technology that upset the new status quo. Figure 5 depicts this emerging conceptual framework – integrating organizational dynamics across stages and transitions matching to related technological requirements.

Stage 1: Initiation

Green's AI story begins with various areas of the organization simultaneously experimenting with AI. Individual employees observe the *"hype with deep learning, ImageNet, and those* *things in the consumer world, we saw those […] as inspiration" (G2)* and recognize the danger of *"IT players entering the hardware space and taking over market share from our divisions; [therefore,] the beginning of AI was also a reaction required from the product-side" (G1)*. With a long history of software development and complex hardware-software offerings, Green's departments do not require an official kicking-off point, but can immediately build on a strong base of engineering and data expertise. Rather than as a coordinated effort, AI research emerges in multiple units "*bottom-up, it was a grassroots-movement. We did have empowerment from the top, but there wasn't a corporate program that was pushed through. It's really surprising, how bottom-up it all started" (G3)*. The strategic goal is clear at this point – securing the firm's market position – but is individually recognized by departments and not communicated centrally.

For Blue and Red, initiation emerges out of similar motives, but more centrally driven. Blue's head of the business intelligence department recognized that *"this [technological trend] is going to be important" (B2)* and approached the CIO to jointly *"build a central hub for analytics in the organization" (B2)*. With individual employees already pursuing their own primitive experiments, these resources and additional hired experts were pulled together to initiate efforts more officially in a team of *"data scientists, data engineers, AI consultants and architects" (B2)*. For these efforts, Blue provided *"central funding, beginning to end" (B2).* Similarly, Red's journey also began with a central push, even formalized in the company strategy communication. Across the organization, leadership recognized the opportunity for change away from slow and tedious processes: *"looking back, it is incredible just how many customized reports there were, that had to be produced regularly. It really was a deeply rooted culture [of manual data processing]" (R2)*. Management responded by setting up a *"central data analytics team very strongly focused on the processes"* as well as a specialized team for the sales platform *"very strongly focused on* *customer interaction” (R1)* – a strongly signaled initiation of AI development, albeit *“there wasn't a clear plan or agenda, it was relatively opportunistic” (R2)*.

Initiation in all organizations is therefore triggered by the growing recognition of developments in the external environment, providing both strategic threats and opportunities. Idiosyncratic preconditions and skillsets then manifest in varying structures along the centralization-decentralization spectrum as a first push towards AI. All organizations pursue similar overarching goals, to stay ahead of the curve of technology development in order to secure profitability through efficiency and market position, but what this means for the organizations and their individual actors’ use of the technology is entirely unclear at this point, as they lack “models of reality […] to generate predictions [and] consider the merit of alternative strategies” (Knudsen & Srikanth, 2014, p. 427).

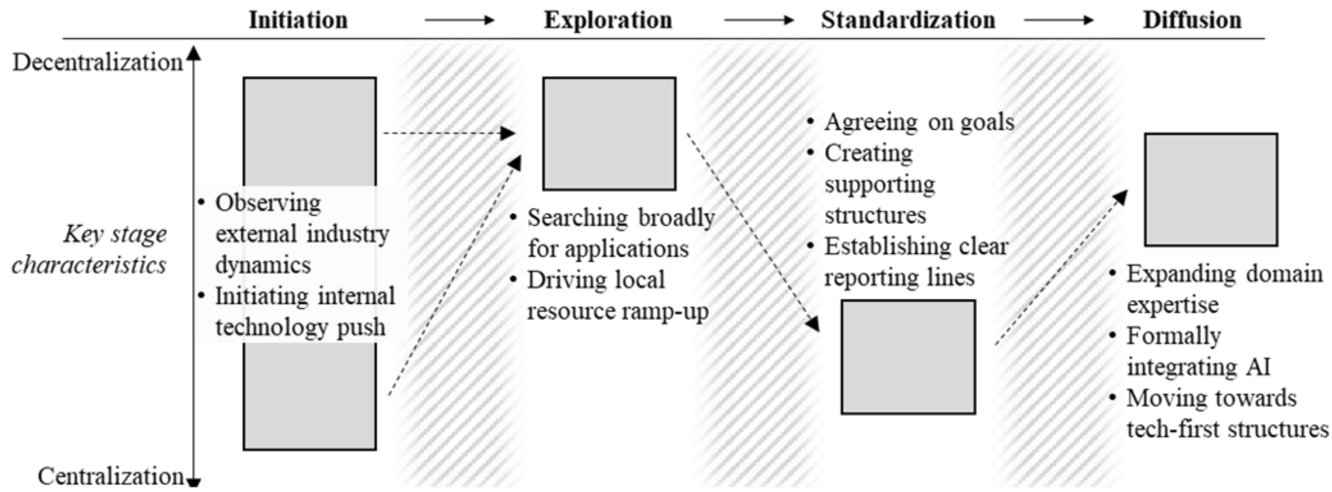

	Transition 1	Transition 2	Transition 3
Type	Creeping development with independent businesses ramp-up	Active political- and performance-driven intervention by HQ	Slow hand-over by HQ *or* business driving domain-specific interests
Typical causes	• Arising organizational limitations • Growing decentral curiosity and opportunity	• Growing unrest with lack of outcomes • Increasing strategic attention • Reshuffling political responsibilities	• Emerging organizational limitations • Accepting loss of control
Technological requirements	• Broad search for use cases • Growing familiarity with tools and opportunities	• Creation of joint tech- and tool-stack • Data availability across silos	• Domain-specification of tools • Ongoing adjustments for business as usual

Figure 5: Emergent process model of GPT introduction

Transition & stage 2: Exploration

Two related factors lead the organizations away from these heterogeneous origins into a similarly broad Exploration phase: firstly, a growing bottom-up curiosity to investigate the most promising use cases for specific domains, which eventually leads to, secondly, the initially designed organizations meeting structural limits. Through these developments, we further recognize the emergence of conflicting goals and differences in underlying mental representations of AI technology between actors.

Red, starting with their initial set-up of two distinct teams in separate organizational units, quickly saw duplications of efforts and investments in both human and technological resources. The teams act disjointly, with *"lack of alignment [and] living alongside each other than with each other" (R2)* in *"an overlapping setup, [each] with their own tech stack and their own approach to prioritizing data science use cases, with unclear roles and responsibilities between the two teams" (R3)*. In the initial set-up, the customer-facing platform team had clear control over this area of application, with the central data team left to cover the rest of the organization and all functional departments. However, other teams now also began experimenting outside of these structures; recognizing the future potential of AI, *"people thought it was cool to create own solutions" (R2)*, and the planned coordination quickly fell apart.

Similar waves of enthusiasm soon swept across Blue, further accelerated by freely available budgets, or slack (Nohria & Gulati, 1996), which meant that *"if you had a fancy data idea, you received required funding fairly easily" (B4)*. Coupled with a lack of coordination driven by personal interests as "*every department head wanted to sharpen their personal AI-profile" (B5)*, departments began individual exploration. Instead of consistently making use of the provided central structures, these individual departments end up in control over the AI initiatives, exploring technological tools, suppliers, and implementation strategies individually and independently. Informants called this phase of search the *"wild west,* *maximum hype" (B5)*, where *"everyone wanted to do something with data, everyone created azure or IBM accounts, the tech companies could sell their stuff separately to every department" (B5)*. The restrictions in capacity, budgetary freedom of departments, and individual interest show clear limitations of the original structures.

On top of these, organizations are also limited by their inability to effectively evaluate the success of attempted AI implementations, which only serves to further uncoordinated exploration. At Green, for example, departments created a *"diversified product portfolio [with] broad interests and use of various technologies" (G1)* and had now *"gone quite far. They had use cases very early and started building up data teams" (G2)*. Many of these projects were not only developing algorithms, but also acquiring new suppliers of tools, coordinating processes, attempting to obtain relevant data from multiple sources, and developing own performance indicators all in parallel. While not aligned across the organization, we see this stage of uncoordinated exploration as a key driver of value in larger organizations. By allowing individual teams, and on aggregate the whole firm, to create variance in the ideas that are being applied, the organization creates a better understanding of the uses and limitations of the technology. However, obtaining and coordinating feedback is difficult under these circumstances, and will often confound multiple aspects of structures. Without reliable predictions stemming from reliable mental representations, firms are left no choice but to first undergo the costly implementation of ideas in order to evaluate them. Uncoordinated exploration leads our organizations into a situation with with "*very many PoCs across the organization - but no tracking of results" (B4)* – out of which a general sense of dissatisfaction soon emerges.

Transition & stage 3: Standardization

Exploration eventually ends in *"disillusionment as projects didn't scale as fast as imagined, producing little impact" (B5)*, for all organizations. They now find themselves in the well-known dilemma of search (Boumgarden, Nickerson, & Zenger, 2012; Fang, Lee, & Schilling, 2010; Siggelkow & Levinthal, 2003); leaving departments to continue individual exploration builds a foundation of interest and expertise, but is costly and *"wasteful in terms of monetary and human resources" (B4)*. As a result, we now see clearly demarcated interventions in all organizations – a shift towards Standardization. These are driven by three factors. Firstly, a growing dissatisfaction with the lack of tangible outcomes, which interestingly leads not only the HQs, but also the affected departments to request more coordinated action. Secondly, an increase in strategic attention as the technology becomes more prominent publicly, and thirdly, shifts in power dynamics through political re-shuffling. Through these interventions, we now also observe how organizations align to form common goals and measures of success, by introducing artifacts and shaping patterns of behavior.

In all organizations, we notice growing general unrest. As projects are *"often stuck in the proof-of-concept phase, lacking department buy-in" (R2),* managers begin *"realizing that we have things going on […but struggling to] really, separate the good from the bad" (G2)*, and eventually look to *"get a measurable return on all this invest" (G2)* by creating central control of proceedings. Of the three case organizations, Red's transition to Standardization was most dynamic, requiring three attempts before reaching stability, each providing clear markers for an increase in strategic attention and political reshuffling. The first attempt begins when both CIO and COO leave the company in short succession, allowing the CFO to integrate analytics and all of IT into their own board division, bringing strategy, implementation, and control together. This centralizes AI-related activities that were previously dispersed, with the noticeable exception of the sales platform team, which maintains its status as key sales channel with independent data analytics. The attempt to centralize is thus also a political struggle between *"powerful individual kingdoms; depending on the internal political situation at the time and considering which kingdom is currently powerful, you want to grab those topics that are en vogue – before things change again" (R2)* – a clear indicator of conflicting goals and misaligned incentives. These political powers had to realign two more times before Red reached a consensus amongst all actors; once shifting even more power towards the sales platform with Covid lockdowns, and a second time with the introduction of a new corporate strategy. Only through this public communication, placing a focus on efficient operations at the forefront of the digital strategy and thereby setting clear goals for the *entire* organization to work towards, did Red manage to implement a system that was stable enough to incorporate all analytics activities under the same umbrella.

Using corporate strategy announcements as a signal, both external and internal, drawing attention to clearly communicated goals is common to all organizations at this stage. As *"the topic became more relevant and interesting strategically" (G2)*, an increase in C-level engagement follows. For Blue, the primary strategic interest was on products; placing AI integration into key product lines front and center of the organization's future through a "*corporate strategy driven by the competitive environment" (B5)*, Blue drew attention to AI development with *"large budgets for product development, which became a lighthouse [for other teams], creating awareness" (B5)*.

With overarching strategic goals now more clearly defined, the organizations continue Standardization by building shared processes and an understanding of potential payoff structures. For Green, this meant merging existing technology platforms, selecting the most advanced *"to serve as base technology stack for both product and process applications" (G1)*. Further, with *"financial resources available, there was this impulse to create a central unit" (G1)* which entailed both a data strategy team and an AI-specific corporate incubator, for *"workshop formats, teambuilding, value proposition design, methods, and use case structure" (G3)*. Eventually this results in an *"inner source library and toolbox, with best practices to re-use for use case types, to approach scaling, become more efficient, become faster, and of course cheaper" (G2)*. All the while, budgetary freedom let *"the AI hype reach its peak" (G2)*, with departments *"knocking down the doors" (G2)* of the experts.

Communication initiatives such as a *"digitization-network for informal exchange" (G4)* and a *"series of events [...] available to all colleagues [...] to give a quick run-down on technical basics" (G2)* mirrored and supported this growing overall engagement, which was monitored through *"a KPI system to track use cases, reported to the global CTO every 6 months" (G1)*. Centralizing development activities thus directly resulted in the organization creating models of AI usage; mental representations that establish a shared vision and path forward, allowing for the simulation and evaluation of applications without costly implementation.

Blue's developments mirrored this, with the top-down *"drive to create clearly specified AI targets per department" (B4)* supervised by a new *"digital strategy unit*" (B4), with the mandate to *"to act outside of IT, centralistic, cross-functional, in touch with the entire organization" (B1)*, also ensuring data access where required. The effects of this measure cannot be overstated; each department, pushed by a central force, was now engaging with AI in a *"a more structured manner – create a playbook and actively manage pipelines of use cases" (B5)*, and importantly also coordinating their efforts across departments and levels of hierarchy. With an economic downturn requiring more prudence, the central units implemented stronger control mechanisms, such as *"decision-boards across multiple levels, aggregate upwards into the board level, using the same tools to document, report, and evaluate use cases" (B5)*. Red's monetary focus was even more prominent. The newly installed digital unit in the COO board division decided to work on only a small number of use cases, selected those with *"very strong buy-in from the respective departments top leadership, only if it directly supports one of the key elements of our strategy… [and] only if the value generated is in the millions, if not billions, over this strategy cycle." (R4)*. Clear tracking, full transparency, and extremely strong top management support across all divisions – an approach vastly different from previous iterations.

These developments make clear how the phase of Standardization in our model differs from the reintegration for the purpose of targeted exploitation we see in literature (March, 1991; Posen & Levinthal, 2012; Siggelkow & Levinthal, 2003). Previously, organizations were seen as using decentralized structures to explore a new field of technology – thereby identifying the most promising avenues to value generation, and subsequently centralizing in order to exploit these same techniques. In our observations, however, organization use centralization as a tool not to exploit value creation, but instead to *create an understanding of value creation* in the first place. We see that Standardization via the creation of processes and communication channels is still part of the process of organizational sensemaking, via the creation of shared mental representations. At this stage, our organizations are still struggling to identify where value lies and how to make use of AI. Driven by unsatisfied expectations from both central strategy units and the departments themselves, we see organizations now come together in a new approach to the creation of mental representations.

While the approach to this Standardization differs slightly between the organizations, all firms thus install structures to fulfil two previously lacking roles; (i) to create a shared understanding, define aligned goals, and subsequently push a broad data strategy and (ii) to support the implementation of use cases across the organization in line with more clearly aligned payoff-structures, tracked through tightly knit processes of feedback and refinement. Formal and informal communication networks, training and mentoring exchanges, KPI tracking systems, implementation toolboxes, and regular reporting and milestone tracking all come together to create a more synchronized and valued system than before. Perhaps even more than for other technologies, this synchronization is important to avoid costly lock-in effects of contradicting path dependencies in the organization. Over time, technology stacks become more advanced, vendors agree to long-term contracts, and data is prepared and stored in specific formats and locations for specific applications. These decisions form a vital part of the understanding of what AI can be – making reintegration of individual solutions across the organization more challenging the longer they are in use, and further increasing the importance on shared representations as generated through coordination during the stage of Standardization.

Transition & stage 4: Diffusion & Platformization

Even with the advantages of a centralized set-up, all organizations recognize the need for an eventual diffusion of responsibilities – and explicitly present this as the goal of their AI strategy. Having now created a shared understanding of the technology, a shared approach on how to address and handle implementation, and created shared goals, key actors in all organizations agree that implementation can only be successful in the long term if the technology is firmly embedded in everyday business processes, building on strong domain expertise and technology-focused structures. This future setup must be decentralized, *"more vertical, with smaller market- and function teams" (R3),* meaning that *"it would be a sign of success to decentralize […] and have data analytics be an elementary part of business as usual" (R4)*. Recognizing the need for the dissipation of skills and the build-up of specific domain expertise, central teams accept that this future may make them redundant: *"the value creation is so essential to them, they want it within their departments […] I hope that AI use will be standard, so that you won't even need us anymore. Maybe dedicated support within the verticals, but not centrally." (G2)*. This would also entail a shift away from traditional hierarchical structures, towards "*a sustainable organization, with decentralized teams, […] in a product-based-organization, [to] change entire organization to become more software-driven, digital, and fast*" for which "*capabilities and network must be decentralized" (B1)*

This differs to suggestions from previous literature, which seems to suggest that a centrally managed organization should be most useful for long-term exploitation of new technologies, to then distribute ideas and developments into the organization as needed (Rivkin & Siggelkow, 2006; Siggelkow & Levinthal, 2003; Smith & Lewis, 2011). However, the specificities of AI require different treatment than previous technologies or even GPTs, in that it requires specific expert knowledge to be implemented efficiently and avoid costly mistakes in interpretation. Whereas access to electricity or the internet can most effectively be managed by central actors for large organizations, AI-specific literature uncovers that algorithmic solutions must be integrated into ongoing processes and hosted in the respective domains for a better understanding of the subject matter. Only this decentralized set-up allows for the correct interpretation of results or appropriate reactions to changing circumstances (e.g., Agrawal, Gans, & Goldfarb, 2020; Brynjolfsson & Mitchell, 2017; Iansiti & Nadella, 2022).

However, rather than diving straight back towards decentralization, letting go of all structure, informants warn of two negative effects of the Diffusion phase. First, *"if everything would be decentralized, we would be re-inventing the wheel five times, trying to solve the same problem, throwing money out the window" (G4)*. Second, as the technology is still evolving and new areas of application may emerge, central financial support may also still be required. Here, informants see it as *"important to provide central investments, [as] businesses are afraid of risks with uncertain returns" (G1)* – HQ teams have more freedom to invest without the need for revenue creation. Rather than pure diffusion of responsibility, the long-term function of the previous phase of Standardization is the creation a platform for technological and functional alignment in the form of a hub-and-spoke organizational model.

The emergence of Platformization and Diffusion predictably comes with the expected levels of organizational friction. As the central organization ramped up expertise and structures during Standardization, departments are now less involved in the officially managed use case process than before. Contrary to this, however, runs the continuously growing general interest in the new technology, which builds up frustration over the implementation and reporting requirements imposed by HQ. The dilemma facing the organizations is now to balance central control, with clear visibility and direction, but limited scope and scalability through artificial bottlenecks in capacity, versus giving control back to the businesses, allowing for more cases to be addresses simultaneously, but reducing the synergies of coordination. Albeit running on similar trajectories and focused on similar long-term technology strategies, the approaches to finding this balance reveal differences between the organizations and particularly the way their central teams handle the loss of influence; either actively managing their loss of control over time, or trying to delay and counteract this loss regardless of counteracting forces in the organization.

We notice these differences as departments begin working on use cases outside of the established structures of control. For some departments in Blue, for example, *"the perspective changed; you now start with your own data instead of waiting for IT"* and *"now we have more expertise, it's the businesses pushing IT, not the other way around. Businesses are always the owner, IT is only enabler" (B5)*. Similarly, in Green, departments make clear that *"we don't have direct contact with the central AI. There is cooperation on the corporate function level. [But] are we also doing our own thing? Yes, absolutely." (G4)* or even that they *"have never even thought of speaking to the central data team. I don't know how they could help us" (G6)*. In Red, there is ambivalence in some of the departments, feeling that "*we do need cross-functional cooperation, to create synergies, but it does always slow us down" (R1)* or *"there is still a strong decentralization-pull; units saying, this is too central, we can't get our own topics through" (R2)*.

For Green and Red, the transition to diffusion thus appears to be cumbersome (and still ongoing at the time of the interviews). Interestingly, for both organizations, the central leadership figures were clearly mandated for AI – and AI only – which lets previously solved power struggles re-emerge. With an eventual transition from Standardization to Diffusion not only bearing risk of missing synergies, but also threatening the positions of power that current central technology leadership holds in the organization, these actors seem more reluctant to give up control. Other departments recognize this hesitance towards change and proceed with the first steps of Diffusion under the radar, without involvement of the central teams. As an example, a representative of Red’s central team states that *“if someone is a data scientist and does data science activity, then they should be in our team in this stage. Nothing else in the rest of the organization. […] You would need decentral data scientists, decentral data engineers, ML ops engineers and so on. That's chaos!” (R4)*. Notwithstanding this strong opinion, other teams have already begun *“integrating teams, shifting them from the analytics units into the business line functions” (R4)*. Similarly for Green, in spite of the clear evidence to the contrary above, the central data strategy team still insists that there was, and should be, transparency *“as far as possible”* and that *“we could be even more stringent in out tracking, without killing innovation, you want more rigor” (G1)*. In these cases, actual power of implementation seems to have slipped away from HQ without their active involvement, decentralization is happening subversively, and the shared ideas of how, why, and where to use AI are breaking apart.

Only Blue seemingly anticipated the potential of emerging struggles and managed to integrate solutions into the organizational structure. With the central team’s task being digitization architecture – that is not just the implementation of AI but long-term organizational change – the team could actively support businesses in accelerating the ramp-up of AI competency without the threat of losing their right to exist. By the time of the interviews, the team felt comfortable saying that they *“stopped pushing machine learning [into businesses], it’s reached the maturity level where I can’t do much more, it’s decentralized now” (B1)* and that they are instead looking for the next big challenges – thereby managing the transition towards Diffusion more actively than other organizations.

As key part of this development, Blue is *“moving towards providing a [technology] platform for the entire company, for big data and advanced analytics” (B2)*, which also entails the clarification of data security concerns, providing tools and APIs, and *“ensuring that we have the required quality of data, making it available where it is needed, through* *modern data architecture […] or visualization tools" (B1)*. Albeit from differing levels of friction in the organizations, the idea of creating a platform as a service for departments to then drive their own use case ideas is common to all. In our second round of interviews, almost two years after the updates on struggles of control, Green now reports on *"dozens of DevOps Teams in different departments" (G3)*, which build and develop solutions independently. At the same time, the informant also reports on an unexpectedly strong interest the departments to *"fill the platform with life, definitely trending in that direction" (G3)*; especially with regard to *"data protection or keeping an eye of compliance, where decentral expertise just isn't as strong" (G3)*. The trend towards platformization, in this case, also includes ongoing conversations in *"the communities we created, to share best practices" (G3)*.

While reminiscent of the previous stages, we note clear differences to Exploration and Standardization; in this platform-driven model, the organization remains more closely knit, makes use of the specifically developed communication channels and processes, but lets departments make key decisions independently. Figure 5 indicates this graphically – with a clear movement towards decentralization, but one that stays shy of the absolute lack of coordination we saw during Exploration.

Transition and stage 5: Renewed Experimentation

At the end of 2022, as we had concluded the original phase of data collection for this project, openAI introduced its ChatGPT tool based on the GPT-3 algorithm. This instantly pushed Large Language Models (LLMs) into the spotlight, making their seemingly unlimited abilities available to the broader public for the first time. Since then, ChatGPT, along with Google's Gemini, Microsoft's Copilot, Meta's Llama, and other tools, has only continued to impress with every new iteration. Anticipating that this seismic shift in AI presence and accessibility may also have affected our case organizations, we returned to the field to discuss potential developments. It quickly became clear that, based on the increase in accountability and competency during Diffusion, departments felt authorized to return to broad exploration. Becoming available at a time of growing autonomy in decentralized structures, this significant shift in AI capability had similar effects to the introduction of an entirely new technology.

Experimentation began immediately, driving the organization back towards decentralization for this new technology cycle (see Figure 6). *"Of course there was another hype moment" (R3)*, and when *"every last person in the business realized what immense effects this could potentially have" (R2)*, many individual employees saw *"how easy it is to set up an account, even with a work e-mail [and began] just playing around – we're a company of developers and researchers" (G3)*. This again led to a situation where departments are doing this *"very independently – it's really difficult to have an overview, even now – but that is part of the culture; there is friction and some loss in efficiency, but it can also be stimulating" (G3)*. Various departments now attempt to *"drive generative AI, as a project" (R3)*, with *"many ideas, a lot of experimentation, hackathons, and so on – but not very structured, at the moment" (R2)*

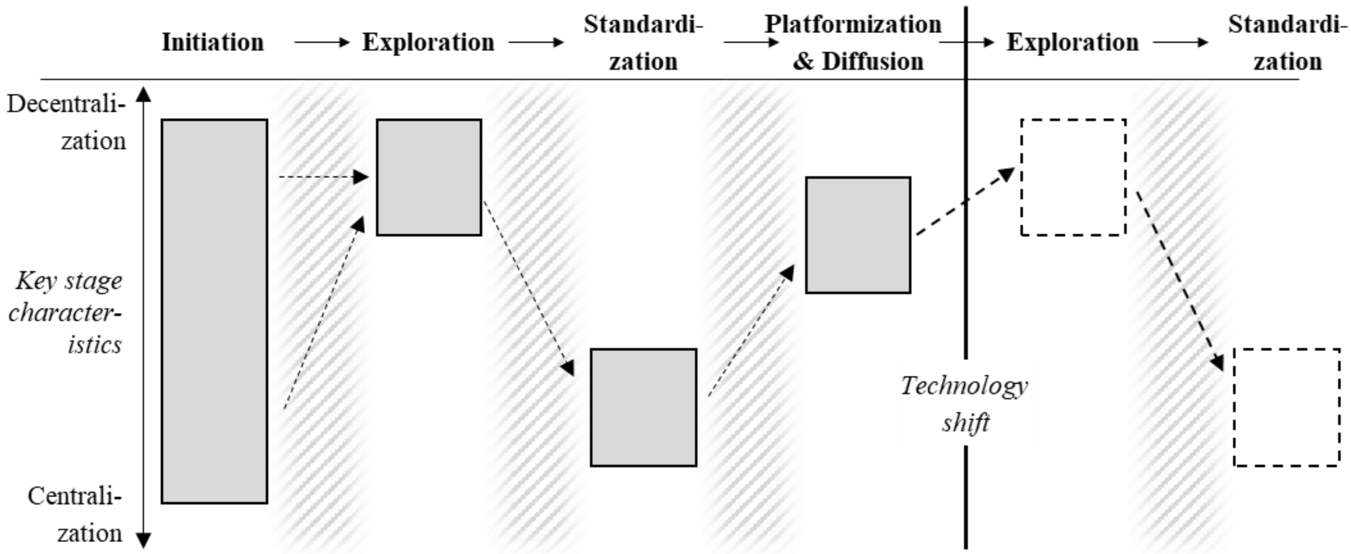

Figure 6: Changing developments after renewed technology shifts

Much like in the original exploration phase, the central organization can only react after the fact; with technology released unexpectedly, available structures are not ready for the increase in uncertainty and the resulting dynamics. The motivation is clear; with a new technology promising to again yield significant value, *"everyone again wants to have a piece of the pie" (R2)*. The advantages, but also the challenges and dangers of this uncoordinated exploration are also clear – and with the central platform of processes and expertise already developed, interventions quickly occur.

Transition and stage 6: Return to Standardization

After the phase of renewed experimentation, central organizations once again react, with the initial response this time triggered primarily by security concerns. Accessing public applications from company hardware and potentially sharing company data with LLMs are, of course, a large security risk – particularly with great uncertainty around how the producers of LLMs store or use data. Again, centralization does not occur to exploit the new technology, as we would expect from previous technological leaps, but instead to understand it, to develop a shared mental representation of how to use the technology to the organizations advantage in the first place. Also reminiscent of the first Standardization phase, we see both sides showing interest in this development. Individual departments realize the lack of direction and threat of dangers as they look to the central organization to provide guidance, as the central organization looks to create synergies, generate value, and mitigate risks.

As decentral departments are *"still trying to understand what the options and limitations are" (R3)*, central teams are now working on *"making sure that we have secure tools, that all users can access quickly. […] It's our job to figure out the happy medium; where we can avoid major risks, but also won't become too slow compared to our competitors" (B1)*. One approach to this is the creation of *"guidelines* – on how to work with *LLMs, that you're allowed to access them, but not share data" (R3)*. More controversial use cases, on the other hand, have to be tested in specifically designed *"playground, where we can work with the technology in a safe environment" (G3)* – continuing the pursuit of experimentation, but with centralized control to facilitate the creation of a shared understanding. Another approach at Green, more forward-looking, is the creation of a technology radar, which *"collects possible genAI use cases that could be interesting – IT and strategy built it together and are now working towards communication to departments" (G3)*. All three approaches show how the case organizations are once again in the process of using the variation created through individual exploration to now create shared mental representations of what this technology can mean for the organization, and the opportunities and threats it creates. *"The pendulum keeps swinging back and forth" (G3)* in terms of the allocation of decision-making authority, but rather than centralizing for exploitation, we once again observe centralization for the co-creation of strategy and goals.

DISCUSSION AND CONTRIBUTIONS

Implications for Theory

The goal of this study is to understand the organizational dynamics affecting the implementation of AI, as a representation of digital general purpose technologies. We present our findings as an emergent conceptual framework of GPT introduction that identifies distinct stages, characterized by alternating centralized and decentralized organizational structures, and distinct periods of transition between them. But instead of being determined by external fit (e.g., Boumgarden, Nickerson, & Zenger, 2012; Siggelkow & Levinthal, 2003), this pattern is the result of internal dynamics and individual attempts at reducing the fundamental uncertainty surrounding AI integration: Organizational structure functions as a catalyst for integrating multiple goals and ambiguous feedback, stressing the importance of structure in high-velocity environments (Davis, Eisenhardt, & Bingham, 2009).

In the first stages of the process, our findings focus on the existence of multiple and conflicting goals. While these are often seen as a hindrance to exploitation (Levinthal & Rerup, 2020; March, 1962; Simon, 1964), we show how the broad exploration of ideas, tools, and infrastructure serves as a positive for organizations in this situation, allowing for the necessary broad exploration that leads to initial mental representations of how AI could serve specific domains. From there, we observe a strict intervention and subsequent coordination through central structures. We show that this is a central element of the creation of shared mental representations; only through coordination can the organizations align and find common ground. A key element at play in this transition are power of decision-making, requiring actors in the dominant coalition to hold sufficient formal and informal recognition. Yet, we also show that the push towards Standardization isn't just centrally driven, but also a result of individual departments requesting support through coordinated approaches. As such, Standardization is not a result of the organizations moving towards exploitation, but rather a necessary step in the creation of shared mental representations.

Our model then shows how central actors are required to let go of the control they have established, leaving departments to complete the integration of tools into their processes with minor support from central platform structures. Finally, we observe a critical change in technological capabilities, which requires organizations to once again re-configure their structures. Building on the now established representations and processes, this reconfiguration takes place with less friction than before. It shows, however, how the diffusion of control over one technological cycle opens door for Exploration with the next. Central control here always lags behind individual interest – and so repeated Standardization may be required again, as a controlling intervention when ideas again begin to threaten the shared representations.

Our findings firstly speak to a recently voiced shortage of research examining the early stages of organizational decision-making, that is the initial information processing to frame goals and problems and its implications on structure (Joseph & Gaba, 2020), instead of later stages focusing on the resulting process of search and evaluation of strategic options. We contribute to this discussion with the creation of our model, addressing this early-stage shortage, and connecting to research that has stressed problem-formulation as central to decision making (Nickerson, Yen, & Mahoney, 2012; Posen et al., 2018) and research focusing on multiple goals in organizations (Ethiraj & Levinthal, 2009; Gaba & Greve, 2019; Levinthal & Rerup, 2020).

This early-stage perspective is a necessary consequence of the conceptual implications of the introduction of a digital GPT. Previous assumptions, which served as the foundation of models discussing dynamic search for organizational structure in the face of past technological change, may no longer apply: Organizations holding shared mental representations of possible future states and their respective payoffs, the ability to evaluate performance feedback, and the ability to easily transition between stages of organizing. As a GPT provides a near endless and unpredictable space of implementation options, organizations must begin searching for possible solutions before understanding the potential payoffs associated with respective choices. Our model therefore serves to provide an understanding of organizational alignment under fundamental uncertainty. By presenting shared representations as the outcome, our findings indicate that organizational goals (or problem-formulation) need not precede a search for solutions but may in fact result from it.

Secondly, our model shows how the creation of mental representations is a function of structural shifts, power dynamics, and individual cognition. While previous research has already provided insights into dynamic structures serving organizations in times of technological change, (Boumgarden, Nickerson, & Zenger, 2012; Knudsen & Srikanth, 2014; Siggelkow & Levinthal, 2003), our model focuses specifically on the periods of transitions between stages of organizing, which serve to shift internal power dynamics through the allocation of decision-making authority. For example, the transition from initial Exploration toward Standardization is crucial in forming shared representations by helping the formation of coalitions (Cyert & March, 1963), which are effective in temporarily resolving tensions resulting from the pursuit of individual incentives. Shifting toward centralization requires key executive positions to be filled with actors of sufficient reach and internal standing – and the interest of decentral departments in making use of more supporting structures. Yet, misalignment between prematurely initiated steps towards centralization and the actual power dynamics between affected parties significantly impedes the effectiveness and speed of structural change.

Our evidence thus points to the role of conflict in organizations (Cyert & March, 1963; March, 1962). Though foundational to the organizational literature, this concept has been left largely unaddressed in current research (notable exceptions aside, e.g., Cronin & Weingart, 2007; Knudsen & Srikanth, 2014), often relying on more benign assumptions about the contexts of innovation (see also, e.g., Joseph & Gaba, 2020; MacAulay, Steen, & Kastelle, 2020).

Finally, instead of exploitation in ultimately central structures (Siggelkow & Levinthal, 2003), we observe a diffusion of responsibility. This lets structures evolve into hybrid designs, which, akin to platform-based structures (e.g., Kretschmer et al., 2022), promise sufficient control while still allowing specialized exploration. Our framework can be interpreted as a process of sensemaking (Cattani et al., 2018), highlighting the socio-cultural dimension of organizational structure (Joseph & Gaba, 2020) and the requirement of paying close attention to the specificities of the technology at hand. As organizations are confronted with uncertainty, structures emerge that first allow individual actors to create their own idiosyncratic understanding, which must then be aligned. Centralization is necessary to facilitate this alignment. Centralization is thus a necessary step for sensemaking, especially in the early-stage development of goals and understanding of payoffs. In this scenario, however, is not the most suitable approach for ultimate exploitation – as all actors ultimately agree, having formed a shared understanding.

Implications for Practice

Building on our findings, we derive a number of strategic recommendations to practitioner audiences involved in the process of introducing AI. We show that AI integration does require careful balance between a hands-on and hands-off approach. Active involvement and intervention are necessary, especially in the movement towards Standardization. However, we also show that too much involvement can be counter-productive, for example when forcing centralization against natural development, or when enforced processes become bottlenecks that lead to mistrust and purposeful avoidance.

Managers therefore face the challenge of having to act counterintuitively for much of the integration process: They must think ahead and prepare the next steps – even as the current situation implies the opposite reactions. Specifically, during Exploration, managers may feel the urge to centralize efforts immediately, putting a halt to the possibly wasteful efforts of individual departments. However, this broad search and familiarization with AI serves to develop general interest and enthusiasm, which we find to be vital to the process. Here, patience is key – but while letting Exploration continue, preparations for Standardization may already be required. When finally attempting Standardization, managers must be aware of the stakeholder landscape, as working against the prevailing interest of strong actors can reset the clock on developments within any given stage. Where leadership may be tempted to force the issue, they must instead wait for – or work towards – the opportune moment where departments are sufficiently engaged, where cross-departmental cooperation becomes a pressing issue, and where there is growing strategic interest from internal and external stakeholders.

Finally, letting go of Standardization once again requires managers to find a delicate balance – this time between continuing to drive the most important use cases through intense resource commitments, and at the same time allowing a maturing decentralized organization to take on the responsibility of developing domain-specific knowledge, building on established tools. Two ideas play an important role at this stage. First, a platformization of structure that allows departments to fall back on established routines, which are continuously managed and maintained by the central organization. Second, the careful framing of responsibility. Creating central departments focused specifically on AI creates a self-preservation interest that works against Diffusion. Creating central departments with a broader scope allows them to hand over responsibility without the fear of redundancy, leading to smoother transitions and more future-oriented thinking.

Future Research Avenues

Our study is, of course, not without limitations. We acknowledge that the interview-based part of our data collection was largely retrospective – which means that we may not have been able to observe such developments and AI use case ideas that were abandoned in the process but potentially still relevant for our theorizing. In turn, having such data could enrich our insights by a performance metric, aimed at deriving recommendations as to what kind of structure would be most facilitative for different types of strategy and use case. While this present analysis therefore focuses on the ease of and disruptions in the implementation process, a further study with an earlier and/or longer observation window could focus on the success of implemented use cases. Further, our analysis includes large multinational corporations which raises questions about the applicability of our findings to smaller organizations, where resources may be limited. Interesting differences may result from smaller organizations as being typically more nimble in transitioning between centralized and decentralized structures, leading to different outcomes.

Similarly, it could be interesting to expand the analysis to a larger sample – perhaps quantitatively – to measure the impact of AI measures on financial indicators such as profit margins. This would require an analysis of a broader market with an in-depth understanding of the efforts of individual organizations measured against their competitors, which could then reveal significant structural implications and the significance of choices made during the process of search. As outlined above, future research could also expand the field of technology introduction towards the process of ‘platformization’ in this context, creating a clearer understanding of the details and requirements for maintaining a functioning intra-organizational platform for a rapidly and dynamically evolving technology.

Chapter 3:
Measuring Organizational AI Integration – A Quantification of Structural Determinants for AI Breadth and Depth

----

Note:
A previous version of this chapter is under review by the Journal of Economics and Management Strategy, at the time of writing. It is also accepted for presentation at the 2024 DRUID conference (Huber, 2024)

The effective use of artificial intelligence (AI) tools in organizations is the subject of a growing field of research, both empirically and conceptually. With AI showing characteristics of a general purpose technology (GPT) and promising to revolutionize not just individual business models, but entire industries (Agrawal, Gans, & Goldfarb, 2022b; Goldfarb, Taska, & Teodoridis, 2023; Iansiti & Lakhani, 2020; Jovanovic & Rousseau, 2005), organizations are increasingly recognizing the need to extract value from algorithms. In keeping with the rise of attention which AI receives in business and public discourse, scholars have begun to provide a deeper understanding of the potential and limitations underlying the technical phenomena, their potential applications, and the mechanisms through which uninitiated organizations may approach AI.

From such previous research, we gain a detailed and valuable understanding of the characteristics of organizations employing AI; organizations using AI are still in the minority, albeit with growing numbers (Rammer, Fernández, & Czarnitzki, 2022). Most large firms, however, already report use, alongside high-growth younger organizations (McElheran et al., 2024). We also understand some of the key drivers for AI implementation, such as a higher value placed on intellectual property rights and process innovation, or having received external venture funding (McElheran et al., 2024). Looking into the processes of those organizations using AI, scholars have investigated the effects of AI use on various measures of performance. AI use can increase creative output on average, for example, but may also reduce the diversity of produced ideas (Dell'Acqua et al., 2023; Zhou & Lee, 2023). Productivity gains may be delayed on an economic scale (Brynjolfsson, Rock, & Syverson, 2019), but can be expected on an individual firm level (Czarnitzki, Fernández, & Rammer, 2023), reliant on complementarities such as high levels of IT capital, flow-efficient processes, or the availability of digitized information (Brynjolfsson, Jin, & McElheran, 2021; Gambardella et al., 2021), thus showing emerging parallels between AI and previous technology-related management practices (Scur et al., 2021).

A significant downside many of these studies share, however, are their limited insights into the prerequisites and process of AI integration – finding suitable structures to enable effective AI use. Instead, we see two patterns: Papers either show AI use as the dependent variable, against which a range of organizational characteristics are measured, or treat AI as a, generally binary or ordinal, independent variable in studies on productivity or innovation. With this, studies are still lacking depth in their analysis of what factors allows organizations to use AI use *effectively*, as many of the complex intra-firm determinants and strategic structural choices remain unobserved. Authors increasingly recognize this as a problem for the field, stating that there is still "restricted availability of data on AI adoption in the business sector" (Czarnitzki, Fernández, & Rammer, 2023, p. 189), that "direct firm-level measures of actual AI use are rare, as are rich organizational 'intangibles'", and that "typically observable firm characteristics leave unexplained a large fraction of the variation in AI adoption across firms […] underscoring the importance of idiosyncratic firm-level unobserved factors" (McElheran et al., 2024, pp. 377, 399). This is particularly problematic in light of the increasing recognition that the capacity to extract value from AI is dependent on suitable organizational structures in reinforcing systems of organizing (Agrawal, Gans, & Goldfarb, 2022b; Bresnahan, 2021; Iansiti & Lakhani, 2020). Considering the increasingly large role AI promises to play, we suggest that the creation of such suitable structures, exploring how to use AI *well*, deserves more attention; an intermediate step between the two existing archetypes of studies.

Using data from a survey designed specifically for this project, we dive deeper into organizational structures, processes, and decision-making frameworks than previous studies. We introduce two new measures to identify not just the use of AI, but the degree of AI integration that results from these decisions: Exploration breadth (measured as the number of ideated use cases) and integration depth (measured as the share of key processes supported by AI tools in a live environment) – and investigate these against a broad range of organizational covariates.

First, we use descriptive tools to compare our data against the results of previous studies on AI *use*, thereby confirming key findings from existing literature. Like previous studies, we find that organizations which are more economically successful, younger, more innovative, and more focused on IP rights are more likely to employ AI, showing statistically significant differences to their non-AI using counterparts (Czarnitzki, Fernández, & Rammer, 2023; McElheran et al., 2024; Rammer, Fernández, & Czarnitzki, 2022). We add as significant driver the association to high-tech industries, but, unlike previous studies, find no significant effects for firm size.

Following calls for an open methodological approach in the systematic empirical investigation of understudied phenomena (Graebner et al., 2023), we then focus on the inner workings of organizations using AI, in order to investigate the drivers of AI exploration breadth and integration depth. We ask whether the identified drivers of AI use also affect the breadth of exploration and depth of integration. Using a series of two-stage Heckman regression models, allowing us to separate selection- from integration-effects, we find that, of the initial drivers of use, only the association to high-tech industries affects exploration breadth (negatively), while only the previous organizational success significantly (and positively) affects integration depth.

Next, we use the same models to identify further organizational correlates of AI integration, zooming in on the departmental- rather than the typically explored organizational level. For AI breadth, we find that a larger invest in digitization measures and stronger decentralization of decision-making both increase exploration. For integration depth, we also see the effect of larger departmental invest in digitization, adding the availability of digital data as significant driver. This matches findings on complementary assets, previously considered on an organizational level (Brynjolfsson, Jin, & McElheran, 2021). We add two further departmental factors as relevant drivers for integration depth: A routine focus on creative- rather than process-oriented tasks, as well as the availability of firm-internal supporting infrastructure. We further attempt to investigate dynamic structures of decision-making. Based on previous literature on technology introduction, (e.g., Burns & Stalker, 1961; Kretschmer & Khashabi, 2020; Miller, 1996; Siggelkow, 2002; Siggelkow & Levinthal, 2003), we use an exploratory margins analysis to relate decision decentralization and organization size to integration depth at various stages of AI use over time.

Our findings primarily contribute to the discourse around the introduction of AI with the creation of new micro-level measures for the degree of AI integration. This reveals a need for more precise language when theorizing around new technology introduction and AI specifically; the factors driving AI use on an organizational level may differ significantly from those driving AI exploration, which again differ from those ultimately driving the depth of AI integration – which may yet be different from those driving ultimate success. Our results therefore make clear that researchers and practitioners must more carefully consider which aspect of AI use they are looking to investigate in particular.

Second, while our findings confirm much of the previous knowledge around AI use and integration, for example by re-emphasizing the value of complementary assets, we add a level of detail by differentiating departmental from organizational factors and focusing on the local allocation of resources. We further add to the conversation a set of previously unstudied, yet significant factors, such as the use of supporting structures for departments introducing AI. Combined with the availability of slack resources through previous business success, such structures could be representative of increased management focus and support – making them a potential focus point for practitioners looking for guidance on successful implementation of algorithmic technologies.

Finally, our analysis opens the field to further research by introducing new measures of decision-making decentralization and providing a first attempt at connecting these to the use of AI technologies. Building on these ideas, we hope that the field may push forward to better understand how AI integration requires dynamic structural responses from organizations over time.

LITERATURE AND MOTIVATION

Observing organizational changes in the wake of technological innovation has a longstanding and established tradition in organization- and management literature. Studying various aspects of structures and performance affected by the introduction of new technology, researchers have, for example, long investigated the adaptation of roles and responsibilities (Barley, 1986; Ichniowski, Shaw, & Prennushi, 1997), structure (Chari & Hopenhayn, 1991; Colfer & Baldwin, 2016; Jaffe, 1986), control of decisions (Barley & Kunda, 1992; Markus & Robey, 1988), or performance (Bresnahan & Trajtenberg, 1995; Brynjolfsson & Hitt, 1996; Robey & Boudreau, 1999).

The rapid pace of development and associated implications of AI, however, pose significant questions regarding the applicability of previous findings. With AI now often classified as a general purpose technology (GPT; Gambardella et al., 2021; Gambardella & McGahan, 2010; Goldfarb, Taska, & Teodoridis, 2023; Jovanovic & Rousseau, 2005), we already know that its downstream effects on organizations may not only impact a broad range of routines or processes (Kawaguchi, 2020; Kim & Upneja, 2021; Sergeeva, Faraj, & Huysman, 2020; Waardenburg, Huysman, & Sergeeva, 2022), but also team structures

(Choudhary et al., 2023; Raisch & Fomina, 2024), creativity and innovation (Amabile, 2020; Bouschery, Blazevic, & Piller, 2023; Dell'Acqua et al., 2023; Jia et al., 2023; Rammer, Fernández, & Czarnitzki, 2022), thereby strategic decision-making (Agrawal, Gans, & Goldfarb, 2022c), performance (Czarnitzki, Fernández, & Rammer, 2023) or even entire business models and industries (Iansiti & Lakhani, 2020). We further know that the introduction of algorithms has the potential to significantly shift power dynamics within organizations (Agrawal, Gans, & Goldfarb, 2019, 2022b), further upending established processes by redistributing decision-making authority.

In literature investigating previous instances of technological shifts, making effective use of an emerging technology is seen as a dynamic process over time, rather than a static one-time implementation of structures. Often, this dynamic nature of exploration is linked to the centralization and decentralization of decision-making authority. Siggelkow and Levinthal (2003), for example, state that organizations should decentralize temporarily to explore, before centralizing authority for effective scaling. Others mirror this notion of adapting through various forms rather than implementing rigid organizational structures, for example through regular adjustments of strategy (Boumgarden, Nickerson, & Zenger, 2012; Clement & Puranam, 2017; Joseph & Gaba, 2020; Rindova & Kotha, 2001; Siggelkow, 2002). With regard to AI specifically, similarly dynamic processes have been conceptually proposed by a number of observers (e.g., Iansiti & Lakhani, 2020; Iansiti & Nadella, 2022), but have not been the focus of empirical investigation. On top of the findings of AI-related measures of use and productivity, we therefore also build on this stream of literature by measuring the (de-)centralization of decision-making structures.

AI adoption

In the abundance of available research findings, the majority seems to focus on either the binary use of AI or the effects of AI on organizational activity, in line with how researchers have investigated the effects of previous technological changes on standard structures of organizations. This reflects significant interest in two questions: *Who uses AI?* and *What are the effects of AI use?* Less considered in this field is the intermediate question of *Who manages to use AI effectively?* Looking at this as an important step of the organizational AI journey, we find that current knowledge is scarce.

We gain first insights by turning to case-based practitioner-literature. Here, authors have pointed out how the introduction of AI is a dynamic process over multiple stages of organizational transformation (Iansiti & Lakhani, 2020; Iansiti & Nadella, 2022; Leonardi, 2020), that it requires readiness and willingness to adjust organizations widely instead of individual point solutions (Agrawal, Gans, & Goldfarb, 2022c; Fountaine, McCarthy, & Saleh, 2019), and that a fundamental management decision choice is to be made between using AI as a tool or as a core element of the future business models (Huber & Alexy, 2024; Kiron & Schrage, 2019). As key requirements for effective AI integration, these findings point to an engagement of top leadership. The willingness to make long-term decisions is key, setting a course for the future with algorithms even at the cost of immediate returns, as well as the readiness of a workforce to engage with a thorough implementation process over an extended period of time. These ideas are reminiscent of the concepts of adaptive and absorptive capacities of organizations; absorptive capacity as "the ability to learn, involves the capacity to recognize new information, assimilate it, and apply it toward new ends… [which] occurs at both the individual and organizational levels", and adaptive capacity referring "to the ability to change… [the] strategic flexibility [that] allows a firm to proact or respond quickly to changing competitive conditions" (Boal & Hooijberg, 2001, p. 517).

Quantitatively, research on the introduction of AI has attempted to measure proxies for some of these factors. Most notably, making use of the US Annual Business Survey, McElheran et al. (2024), show how factors such as an organizations' funding by venture capital, early business success, and even the owners motivation in founding a business can impact an organizations' readiness to implement AI. They further find that factors representing innovation and innovative tendencies, such as the valuation of intellectual property rights, are key indicators of AI use – matching the ideas of absorptive and adaptive capacities. From this valuable investigation, we not only know some of the key drivers of AI use, but also that the overall adoption rate of AI in organizations was only at 5.8% in 2017, albeit increasing to over 18% when weighed for employment, indicating that most large organizations already used AI. Nonetheless, these and other authors point out that there is "restricted availability of data on AI adoption in the business sector" (Czarnitzki, Fernández, & Rammer, 2023, p. 189) and that other "typically observable firm characteristics leave unexplained a large fraction of the variation in AI adoption across firms" (McElheran et al., 2024, p. 399).

AI integration and organizational correlates

Summarizing the above, we understand which organizations have begun using AI and what their key characteristics are. We also qualitatively understand a good process of AI integration, and how AI might have a significant impact on an organizations' future. The missing puzzle piece, however, is to understand which organizations might achieve this success more easily than others; which organizations may be best able to make use of the AI tools they are attempting to integrate, and what structures are best suited to this process.

The most obvious dependent variable to use in this investigation would be the financial impact of AI use. This is, however, hampered by a significant challenge. We are currently in what we call The Between Times of AI; the time "between the demonstration of the technology's capability and realization of its promise reflected in widespread adoption" (Agrawal, Gans, & Goldfarb, 2022b, p. 4) – before AI unveils its full impact on profits and costs (Brynjolfsson, Rock, & Syverson, 2019).

Instead of measuring financial success, we therefore propose measuring the two key aspects of technology integration; exploration and exploitation (Benner & Tushman, 2003; March, 1991; Siggelkow & Levinthal, 2003). These measures are more direct than financial success, as uninhibited by market, customer, or competitor reactions, and directly affected by managerial and organizational choices. Measuring exploration as the breadth of search and exploitation as the depth of AI integration, we therefore build not only on previous literature on organizing around new technologies, but also on AI-related theory showing how deep and system-level integration is a direct prerequisite for ultimate organizational success (Agrawal, Gans, & Goldfarb, 2022b; Iansiti & Lakhani, 2020). With these dependent variables, we chose a set of organizational correlates that reflect the obvious and less obvious dependencies that may underly the degree of AI use – the "idiosyncratic firm-level [and previously] unobserved factors" (McElheran et al., 2024, p. 399).

DATA AND METHODS

Data collection

The goal of this project is to identify structural factors influencing AI integration in organizations. Since no datasets investigating these factors are readily available, we collect proprietary original data using a survey amongst organizations in a major European economy. Making use of an existing panel with over 350.000 registered users as a representative sample of the population, accessed via a 3rd party market research institute, the survey is conducted using computer-aided web interviews (CAWI). Members of the full panel are first filtered to identify employees (including owners and leadership) of organizations with more than 250 employees, thus ensuring sufficient size of the organizations to identify structural peculiarities and filtering out non-relevant groups such as students, pensioners, or unemployed respondents. Randomly selected members of the panel then receive an invitation to participate in the online survey – with a total of 1049 respondents entering the survey. From our survey participation statistics, we find that of these participants fitting the requirements, 331 reported actively using or being aware of AI use in their organization. This implies a 31.6% AI usage rate for organizations larger than 250 employees. While this number is significantly higher than recently published results across all organizations from the US Census Survey

(McElheran et al., 2024), it is in line with other studies focusing on larger organizations (IBM, 2022b), providing initial validity for our dataset.

All respondents then complete the first part of the survey, consisting of general questions on the industry, organization size, age, success, and innovation history. To identify respondents who can provide accurate information on the in-depth characteristics of AI use required for the analyses described, we first provide respondents with a purposely broad and general definition of AI tools, based on previously conducted innovation surveys (Rammer, Fernández, & Czarnitzki, 2022): “For the purposes of this survey, AI describes tools of data processing used by computer algorithms for independent problem solving. This may include, for example, tools for language recognition, image recognition, or other machine learning applications.” We aske three specific screening questions:

1. Does your company employ artificial intelligence tools?
2. Does your team/ department use artificial intelligence tools?
3. Please specify to what extent you are able to comment on the use of artificial intelligence tools in your organization or department

Reponses of “*No AI use in the organization*” for question 1 were filtered out, as were “*I don’t know*” for questions 1 or 2, and responses of “*I can not provide any information*” or “*I can only provide limited information*” for question 3. Through this process, we identified those respondents who not only use AI in their organizations but are also sufficiently deeply involved in the process of AI introduction to provide the accurate and in-depth information required for our analysis. Using this approach, we collect a total dataset of 393 responses – of which 331 were using AI, 62 not using AI. With further gaps in the data, however, the usable dataset is reduced to 264 responses total, with 230 AI-using organizations and 34 non-users. The final dataset of n=264 is not a representative sample of the entire population, especially regarding the share of AI-using organizations. We therefore do not conduct in-depth analyses on the types of organizations using AI in this study, but instead focus on those analyses which arise out of the models sufficiently supported by statistically significant findings in this dataset.

Similarly, the dangers of collecting independent and dependent variables in the same survey are clear, as are the dangers of letting individual respondents shape the observed outcomes for an entire organization. Working around these limitations, we make sure to focus our analysis on structural factors of organizations that are as close to objective truths as possible. While still being limited by respondents possibly not knowing the facts, this ensures that biases of personal impressions and recollection are avoided as best as possible.

Defining dependent variables

The goal of this investigation is to identify the key determinants of organizational AI integration, as shown by the study design in Figure 7. As such, the key contribution and challenge of this project is to identify survey measures adequately representing AI integration – focusing on the depth of integration through process implementation and breadth of integration through use case ideation. To represent depth, we included two items in our survey: the intensity of AI use in the department over the previous year as measured on a 5-step ordinal scale [*not used / tested AI tools, but no live use / low usage / medium usage / regular usage*] (based on the U.S. Annual Business Survey, as used in McElheran et al., 2024) and the share of key departmental processes that are supported by AI tools [*as a percentage of all processes*]. Checking for measure consistency of these two items using a Cronbach's alpha test resulted in a score of 0.594, indicating a sufficient level similarity to allow for re-combination using a principal component factorization. Using the *principal* command from R's *psych* package, we identified an explained variance of 0.736 of the two constituent input variables and thus created the final dependent variable depth of AI integration in key live processes. The variable was normalized to a scale of 0 to 100, meaning that each unit increase can be roughly interpreted as a percentage increase in integration depth, to be analyzed with GLM regressions.

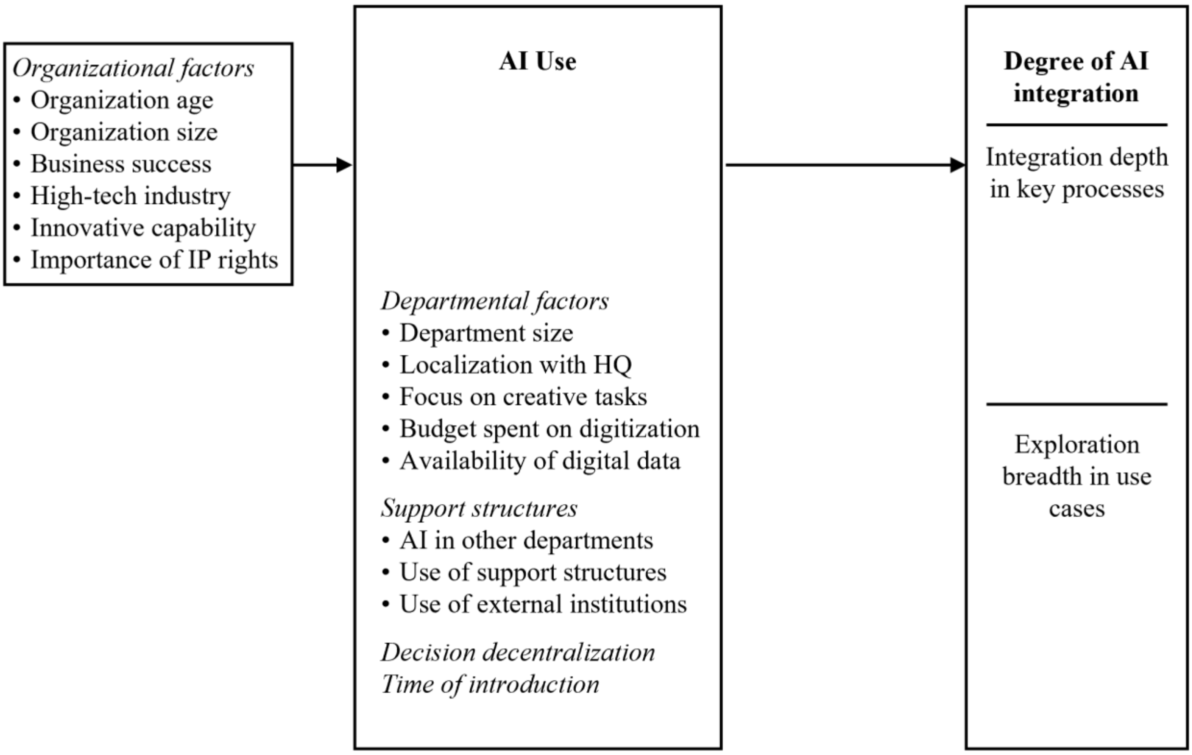

Figure 7: Study design – effects of organizational, departmental, and supporting structures on the degree AI integration

The measure of AI exploration breath, representing use case ideation, is more straightforward: Respondents were asked how many use cases their department is currently working on. With this being a discrete variable, we used a Poisson distribution for all models on integration breadth, first calculating the stage 1 regression and resulting Inverse Mills Ratio (IMR) and using this as input for the stage 2 regression. A key issue in identifying exploration breadth is at what stage of the ideation funnel to make the observation. Our goal is to create a measure of ideation that allows for a fair comparison between organizations. Observing too early in the funnel risks the respondent not being aware of the majority of ideas, dependent on their role, resulting in a measure that is more reliant on the position of the respondent than the actual ideation of the organization. Making the observation too late, on the other hand, by asking for successful or important use cases implemented long-term, reduces the impact actual idea generation has on the measure. We therefore chose the approach of asking for the number of use cases currently worked on by the department, not specifying whether these are employed in live settings or in testing environments – a stage early enough in the funnel to adequately represent idea generation and experimentation, while also late enough to ensure that all respondents, regardless of their role in the department, should be aware of the use cases. While sacrificing some accuracy in the measure of ideation, this ensures comparability between all respondent organizations.

Table 9: Definition of regression variables and descriptive statistics

Variable Type	Variable Name	Description	n	Median	Standard Dev.	Mini-mum	Maxi-mum
DVs	Use	Indicator of any AI use in the organization [*binary*]	264	1	0.34	0	1
	AI_ usecases	Breadth of AI experimentation: Number of use cases worked on by the department [*interval 0-x*]	224	6	2.91	1	25
	AI_ integration	Depth of AI use: PCF variable, combined share of key department processes supported by AI and perceived intensity of use [*interval 0-100*]	230	58.17	20.14	0	100
Organization descriptors	Industry_HT	Industry classifier as high-tech [*binary*]	264	1	0.49	0	1
	Orgsize	Size of organization; number of employees by categories [*ordinal 1-8*]	264	5	1.68	3	8
	Orgage	Age of organization; years by categories [*ordinal 1-5*]	264	3	0.89	2	5
	Orgsuc	Perceived business success against expectations over last 3 years [*ordinal 1-5*]	264	4	1.08	1	5
Innovation-focus	Innovation	Classifier of innovation-focus; new products or processes in last 3 years [*binary*]	264	1	0.25	0	1
	IPrights	Importance of IP rights [*ordinal 0-3*]	264	3	0.59	0	3
Department descriptors	Deptlocal	Focal department co-located at org. HQ [*binary*]	230	1	0.33	0	1
	Deptsize	Size of focal department by categories [*ordinal 1-6*]	230	4	1.28	1	6
Department AI enablers	Deptcreat	Share of creative tasks in department [*interval 0-100*]	230	40	21.02	1	100
	Deptdigi	Share of department budget invested in digitization measures over last 3 years [*interval 0-100*]	230	35	18.85	1	85
	Deptdata	Availability of data in digital form [*ordinal 1-5*]	230	4	0.70	2	5
AI introduction process	AI_otherdept	AI use in other departments [*ordinal 1-5*]	230	4	0.96	1	5
	Time_org	Duration of AI use in the organization [*ordinal 1-5*]	230	3	0.91	1	5
	Support	Any supporting institutions available for the focal department [*binary*]	230	1	0.24	0	1
	S_extern	External supporting institutions available [*binary*]	230	1	0.48	0	1
	DecDistance	Hierarchical distance between AI decision-makers and board of management [*interval 1-x*]	230	2	1.33	0	6
	DecBoard	Indicator of board involvement in AI decisions [*binary*]	230	1	0.37	1	2

Measuring correlates

In the following, we describe the identification and definition of all independent variables used in the regression analysis – the correlates of AI integration. Table 9 shows an overview of the variables and descriptive statistics[9].

Organizational factors

The first category of correlates are characteristics and structural determinants of the organization. These contain the most straightforward identifiers such as organizational age and size, both as ordinal variables. The categories for organizational size are set as *fewer than 100 employees, 100-250 employees, 251-500 employees, 501 to 1,000 employees, 1,001 to 2,500 employees, 2,501 to 5,000 employees, 5,001 to 10,000 employees, and more than 10,000 employees*. For age, the categories were set as *less than 3 years old, 3 to 10 years, 11 to 50 years, 51 to 100 years, and more than 100 years*.

Other items in the category of organizational descriptors rely more on the respondents' interpretation of the organizational culture and performance. These include business success, where respondents were asked to rate the business performance over the previous three years as measured against the organizational expectations *[far worse than expected – slightly worse than expected – as expected – slightly better than expected – far better than expected].* Rather than ask for absolute revenue figures, which could be used to corroborate organizational size and growth (as done for example in the Annual Business Survey, McElheran et al., 2024; Zolas et al., 2020), we choose to introduce a measure of business success against performance targets, or aspiration levels, know to trigger search for performance improvements (e.g., Cyert & March, 1963).

Similar to other studies (McElheran et al., 2024), we ask respondents to rate the importance of Intellectual Property (IP) and its protection for the organization *[not important*

[9] Table 19 in the appendix shows the correlation matrix for all used variables

at all – rather unimportant – rather important – very important] – indicating the organization's treatment of IP as a proxy for an innovation-focus. An additional, corroborating measure for the *innovation-focus of an organization* comes from the straightforward question of whether or not the organization has, over the past three years, introduced new products, services, or processes that were significantly different to competitors' products or services and previously employed processes. These insights were split over two questions in the survey, covering new products/ service innovation in one question and processes innovation in the other. We found, however, that the two constructs were strongly related, with a pairwise correlation of 0.46 and Cronbach's alpha of 0.604, and thus combined both into one innovation variable: the overall, binary characteristic as a previously innovative organization.

Finally, we considered the organization's overall competitive environment and its technical propensity by indicating high-tech industries. Respondents were asked to select their industry out of a list of 28 options or add their own – which were grouped into high-tech or low-tech industries by combining the respective classifiers used by the U.S. Department of Commerce and the OECD (Department of Commerce, 2017; OECD, 2015).

Departmental factors

On the departmental level, we once again use size as a category-based control variable, alongside asking whether the department is physically in the same geographic location as the headquarters of the organization. Based on survey items from literature (Brynjolfsson, Jin, & McElheran, 2021; Scur et al., 2021), headquarter co-location could imply better access to top-level decision-makers, but also mean less flexibility and freedom in decision-making for the department.

The three remaining items on the departmental level entail AI enablers, or complementary assets, that may support the successful integration of new technologies into an organization (Aral, Brynjolfsson, & Wu, 2012; Brynjolfsson, Rock, & Syverson, 2019;

Choudhury, Starr, & Agarwal, 2020); a departmental focus on creative tasks, the budget spent on digitization initiatives, and the availability of data in digital formats. The first of these is based on existing surveys asking organizations to classify the type of activity most prevalent in their operational processes, as done, for example, in the U.S. Census Management and Operations Practices Survey (Bloom et al., 2016). By asking for the share of regular department processes that engage with creative problem-solving such as innovation, this item gives us a good understanding of the type of work the respondent department covers. Notably, this is different from asking for the company division the department may belong to – even in more creative divisions such as R&D, individual departments may be very process-focused. On the contrary, in more process-oriented divisions such as sales, individual departments may be focused on creative marketing techniques. Focusing on the departmental share of creative processes therefore gives us more precise information than asking for the corresponding division.

We identify the share of available departmental budget spent on digitization initiatives using an open percentage field, while the availability of data in digital formats is entered as categories of *very poor data availability – rather poor availability – medium degree of availability – good availability – all required data available*, again based on Bloom et al. (2016).

Support structures

Building on the idea of complementarities on the departmental level, we introduce the idea of measuring the use of supporting structures to aid in the integration of AI tools. We consider three possible features as supporting structures: First, the use of AI in other departments (as known to the survey respondent), enabling both formal and informal exchanges between colleagues. Second, targeted and dedicated internal support structures such as a center of excellence, incubators, or internal consultancies, both as free and paid services. Third, dedicated external structures such as third-party consultancies or state-funded digitization programs. Each of these options was separately asked in the survey, but ultimately combined during analyses to form binary responses for “used any type of support structure” and “used external support structures”.

Time of AI use and (de-)centralization

The last set of variables in the survey relate to the previously mentioned literature around the decentralization of decision-making authority and dynamic adjustments processes over time. For this final analysis, we measure the degree of decentralization using the hierarchical distance between the board of management and the organizational level where AI use cases are identified and selected. Intuitively, this number is lower in organizations where the board of management is closely involved in key decisions around AI and higher in organizations where departments have more freedom to decide their own AI strategies, possibly leading to a greater diversity of ideas within the organization (Boumgarden, Nickerson, & Zenger, 2012; Koçak, Levinthal, & Puranam, 2022; Posen & Levinthal, 2012; Siggelkow & Levinthal, 2003). We then create interaction effects between this decentralization distance and organization size, as well as between these variables and the time of AI use in the organization. The intention here is to reveal how, depending on size, increasing experience with AI may alter the effects of decentralization.

Model overview and descriptive statistics

We create a series of regression models and model variations for analysis (see Table 10 for an overview). For every model, we create 10 variations of modified variables and interactions. The first version is the baseline, including all organizational, departmental, and support variables, without interactions. Variation 2 introduces the first interaction term between organization size and decentralization distance to identify how the interplay of these affects AI integration success, as described in the previous section. Variation 3 introduces the time of AI use for the exploratory analysis, with variations 4 to 10 showing the possible interaction effects between organization size, decentralization of decision-making, and the time of AI use, including a possible U-shaped effect through the squared term of time of use.

Table 10: Overview of model specifications and variations

Model	**Description/ variation**	**1** All controls, incl. Size, Decentral.	**2** Adding interaction: Size*Dec.	**3** Including Time of AI use, no interactions	**4 / 5 / 6** Adding interactions: Size * Decentral. * Time	**7** Including Time^2, no interactions	**8 / 9 / 10** Adding interactions: Size * Decentral.* Time^2
B	Model on **AI integration breadth**, number of use cases	*B1*	*B2*	*B3*	*B4/5/6*	*B7*	*B8/9/10*
D	Model on **AI integration depth**	*D1*	*D2*	*D3*	*D4/5/6*	*D7*	*D8/9/10*
				Post-hoc corroborating models			
I	Adjusted measure of decentrali-zation as **Board-level involvement**	*I1*	*I2*	*I3*	*I4/5/6*	*I7*	*I8/9/10*
R	Introducing **departmental AI readiness** as new PCF variable	*R1*	*R2*	*R3*	*R4/5/6*	*R7*	*R8/9/10*
P	Adjusted to only use **AI process coverage** as dependent variable	*P1*	*P2*	*P3*	*P4/5/6*	*P7*	*P8/9/10*
Y	Adjusted to only use **AI usage intensity** as dependent variable	*Y1*	*Y2*	*Y3*	*Y4/5/6*	*Y7*	*Y8/9/10*

All models and variations make use of a Heckman regression analysis to account for possible sample-induced endogeneity (Certo et al., 2016). As described above, our goal is *not* to discuss the usage of AI in general – but to identify which structural factors affect the integration success of AI once an organization *does* use it. From the original sample, we therefore aim to first control for AI use – and then conduct the analysis from the reduced second sample. The first stage of the Heckman models predicts the use of AI in the organization as a binary variable, including the independent variables of high-tech industry, organization size, age, success, past innovation, and IP importance. The Heckman exclusion restriction used is the variable organization age, significant for the prediction on AI use in all models. The second stage of the model includes all available independent and control variables, modeled against the dependent variable of AI integration depth and use case breadth in models D and B, respectively.

Table 11: Descriptive analysis of key organizational factors for AI-using and non-using organizations

Variable	Average of AI non-users	Average of AI users	Difference	t-test (p-value)
High-tech industry [0/1]	12%	67%	0.55***	- 8.61 (0.0000)
Organization size [1-5]	5.35	5.42	0.07	- 0.22 (0.8236)
Organization age [1-5]	3.94	3.25	-0.69***	4.15 (0.0002)
Org. success in previous years [1-5]	3.12	3.73	0.61***	- 3.11 (0.0033)
Innovation [0/1]	65%	98%	0.33***	- 3.95 (0.0004)
Importance of IP rights [1-4]	2.06	2.72	0.66***	- 4.44 (0.0001)

RESULTS: DRIVERS OF AI USE AND INTEGRATION

Indicators of AI use

We begin our investigation with a descriptive analysis on the effects of organizational characteristics on AI use. Table 11 shows the average values of AI-using and non-AI-using organizations for six key organizational characteristics, building on existing knowledge from literature: the association with a high-tech industry, the size of the organization, its age, its business success in previous years, its track record in innovation, and its value placed on IP rights. Like previous studies, we find a significant positive difference between AI users and non-users in terms of their previous organizational success (Czarnitzki, Fernández, & Rammer, 2023), as well as their interest in innovation (Rammer, Fernández, & Czarnitzki, 2022) and IP rights (McElheran et al., 2024). Comparing the measured innovation tendencies to Rammer, Fernández and Czarnitzki (2022), however, who found an 8% and 8.5%

difference for organization driving process and product innovation respectively, our findings show an even larger difference, with a 33% increase for organizations using AI and driving any type of innovation, as illustrated in Figure 8.

Figure 8: Descriptive analysis of organizations using AI vs. not using AI

We further find that AI-using organizations are significantly younger than non-using organizations. This is line with recent findings by McElheran et al. (2024), who also show that although larger organizations are most likely to use AI, when controlling for size, age is a negative contributor to the likelihood of AI use. Unlike in this study, however, size is not a significant contributor in our sample. Lastly, working in high-tech industries appears to significantly affect AI-use on an organization level, increasing the likelihood of AI use by a staggering 55%, confirming the intuitive belief that these industries are not only working *on* technologically advanced products and services, but they are also working *with* such technologies in their internal processes.

AI Use vs. Exploration Breadth

Diving deeper into the analysis, we now compare the drivers of AI use to those of AI exploration breadth and integration depth. Table 12 shows the Heckman regression output for selected versions of these models.

We first focus on variation B1; measuring exploration breadth, in the static model with all variables, before the introduction of any time-related effects. The upper part of Table 12 (that is, looking at the likelihood of using AI as the first stage of the Heckman model) confirms the findings from our descriptive analysis above[10]: That high-tech, younger organizations are more likely to use AI, as are organizations more interested in innovation and IP protection (all to a 1% significance-level), along with more successful organizations (albeit only to a 10% significance level). Similarly, we cannot confirm an effect of firm size on AI use, with the positive indicator showing no statistical significance.

Table 12: Correlates of AI usage and AI integration, for selected model variations
Key findings and differences in bold

n=264 for all models	B1	B2	B6	B10	D1	D2	D6	D10
DV: AI Use								
Intercept	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-2.99*** (1.04)	-2.99*** (1.04)	-3.02*** (1.02)	-2.99*** (1.03)
High-tech industry [0/1]	**1.31*** (0.34)**	**1.31*** (0.34)**	1.31*** (0.34)	1.31*** (0.34)	**1.31*** (0.33)**	**1.31*** (0.33)**	1.31*** (0.33)	1.31*** (0.33)
Organization size [1-5]	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)	0.1 (0.09)	0.1 (0.09)	0.1 (0.09)
Organization age [1-5]	**-0.39*** (0.13)**	**-0.39*** (0.13)**	-0.39*** (0.13)	-0.39*** (0.13)	**-0.4*** (0.13)**	**-0.4*** (0.13)**	-0.4*** (0.13)	-0.4*** (0.13)
Org. success in previous years [1-5]	**0.23* (0.13)**	**0.23* (0.13)**	0.23* (0.13)	0.23* (0.13)	**0.23* (0.12)**	**0.23* (0.12)**	0.23* (0.12)	0.23* (0.12)
Innovation in products or processes [0/1]	**2.02*** (0.36)**	**2.02*** (0.36)**	2.02*** (0.36)	2.02*** (0.36)	**2.01*** (0.37)**	**2.02*** (0.37)**	2.02*** (0.37)	2.02*** (0.36)
Importance of IP rights [1-4]	**0.8*** (0.2)**	**0.8*** (0.2)**	0.8*** (0.2)	0.8*** (0.2)	**0.8*** (0.2)**	**0.8*** (0.2)**	0.8*** (0.2)	0.8*** (0.2)

[10] The other models and model variations, in this table and in the appendix, show very similar or identical results for the regression on AI use. Minor differences likely occurring due to variations of “NA” values in the respective modified variables of each model

n=264 for all models	B1	B2	B6	B10	D1	D2	D6	D10
DV: AI integration	***Exploration breadth***				***Integration depth***			
Intercept	1.00**	1.16**	0.39	-1.32	-35.7***	-29.88***	-15.33	61.18
	(0.44)	(0.47)	(0.79)	(2.25)	(9.63)	(10.63)	(20.02)	(56.69)
High-tech industry [0/1]	-0.24**	-0.24***	-0.24***	-0.23***	-0.6	-0.83	-0.92	-1.11
	(0.1)	(0.09)	(0.09)	(0.09)	(2.25)	(2.25)	(2.18)	(2.24)
Organization size [1-5]	-0.03	-0.05	0.07	0.49	-0.12	-1.05	-4.68	-14.31
	(0.02)	(0.04)	(0.12)	(0.41)	(0.62)	(0.92)	(3)	(9.09)
Org. success in previous years [1-5]	0.01	0.01	0.03	0.03	2.27**	2.43**	1.96**	2.27**
	(0.03)	(0.03)	(0.03)	(0.03)	(0.96)	(0.97)	(0.98)	(1.00)
Innovation in products or processes [0/1]	-0.16	-0.17	-0.15	-0.14	-3.13	-3.35	-2.63	-2.46
	(0.32)	(0.31)	(0.31)	(0.29)	(3.84)	(4.01)	(3.98)	(4.25)
Importance of IP rights [1-4]	-0.02	-0.02	-0.01	0.01	2.94	2.88	2.64	2.32
	(0.08)	(0.08)	(0.08)	(0.08)	(2.2)	(2.21)	(2.17)	(2.19)
Department located at HQ [0/1]	-0.03	-0.02	-0.05	-0.02	3.32	3.63	4.67	4.88*
	(0.12)	(0.13)	(0.12)	(0.1)	(2.9)	(2.92)	(2.99)	(2.87)
Dep't size [1-5]	0.02	0.02	0.03	0.04	1.1	1.17	0.89	0.73
	(0.03)	(0.03)	(0.03)	(0.03)	(0.88)	(0.88)	(0.9)	(0.87)
Dep't focus on creative tasks [% of tasks]	0.002	0.002	0.002	0.002	0.28***	0.27***	0.25***	0.25***
	(0.002	(0.002	(0.002)	(0.002	(0.05)	(0.05)	(0.05)	(0.05)
Dep't digi.-invest [% of budget]	0.005*	0.005*	0.003	0.003	0.24***	0.23***	0.26***	0.26***
	(0.003)	(0.003)	(0.002)	(0.002)	(0.06)	(0.06)	(0.06)	(0.06)
Dep't availability of digital data [1-5]	0.03	0.03	0.02	0.01	9.02***	8.95***	9.09***	8.64***
	(0.06)	(0.06)	(0.06)	(0.06)	(1.66)	(1.67)	(1.66)	(1.72)
AI usage in other departments [1-5]	0.03	0.03	0.03	0.03	-0.19	-0.17	-0.48	-0.47
	(0.04)	(0.04)	(0.04)	(0.04)	(1.03)	(1.04)	(1.05)	(1.04)
Support structures used [0/1]	-0.08	-0.08	-0.04	0.05	16.96***	16.99***	17.03***	15.32***
	(0.24)	(0.24)	(0.18)	(0.15)	(5.59)	(5.54)	(5.53)	(5.36)
External support structures used [0/1]	-0.02	-0.03	0.02	0.01	2.21	2.13	1.45	1.73
	(0.06)	(0.06)	(0.07)	(0.07)	(1.96)	(1.96)	(2)	(1.96)
Decentralization: Distance to board [0-x]	0.07***	-0.01	0.65**	1.86**	-0.52	-3.47	-10.59	-50.36**
	(0.02)	(0.09)	(0.32)	(0.85)	(0.62)	(2.31)	(8.71)	(22.77)
Time of AI usage in organization [1-5]			0.17	1.34			-2.9	-55.28
			(0.25)	(1.53)			(5.13)	(36.26)
(Time of AI use)^2				-0.18				8.35
				(0.24)				(5.44)
Org.Size x DecDistance		0.01	-0.1*	-0.34**		0.5	2.18	7.9**
		(0.01)	(0.05)	(0.15)		(0.37)	(1.5)	(3.92)
Org.Size x Time of use			-0.03	-0.32			1.03	7.95
			(0.05)	(0.29)			(0.84)	(5.7)
Dec.Distance x Time of use			-0.2**	-1.09**			1.92	31.17**
			(0.1)	(0.56)			(2.5)	(14.46)
Org.size x Time^2				0.05				-1.1
				(0.05)				(0.84)
Dec.Distance x Time^2				0.15*				-4.85**
				(0.08)				(2.21)
Org.Size x Dec. Distance x Time			0.03**	0.2**			-0.47	-4.67**
			(0.02)	(0.1)			(0.41)	(2.36)
Org.Size x Dec. Distance x Time^2				-0.03*				0.69**
				(0.01)				(0.35)
Error/ regression definition terms								
sigma					13.01***	12.97***	12.81***	12.69***
					(0.6)	(0.59)	(0.58)	(0.57)
rho					-0.2	-0.22	-0.2	-0.24
					(0.28)	(0.3)	(0.27)	(0.31)
IMR	0.18*	0.18*	0.16*	0.15*				
	(0.09)	(0.09)	(0.09)	(0.09)				
Adjusted R-squared	0.1037	0.1028	0.1301	**0.1583**	0.5544	0.5556	0.5573	**0.5589**

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

In the bottom half of model B1 in Table 12, the dependent regression variable changes to be our measure of AI exploration breadth – the number of use cases pursued in the respondent's department. Comparing the effects of the same independent variables, now measured against AI breadth rather than AI use, reveals interesting findings. The first clear difference is that almost all variables addressed in the analysis of AI use lose their significance for exploration. Previous organizational success, innovation history, and importance of IP rights all show no effect on the breadth of AI exploration in our data[11]. This is particularly surprising for previous success, when we consider success to be indicative of the resources available to the organization, usually conceptualized as organizational slack. Slack is often considered to expand an organization's perspective on innovation, allowing it to search more broadly for riskier applications (George, 2005; Greve, 2003; Nohria & Gulati, 1996). As such, we would have expected success to increase exploration.

The only significant variable from the regression on AI *use* that remains significant for *breadth* is the association to a high-tech industry – however with the positive sign switching to negative. This means that departments from high-tech organizations identify or explore 0.25 *fewer* use cases than similar departments from lower-tech industries. While having a high-tech background thus makes it more likely that an organization or department begins working with AI in the first place, it also appears to narrow their focus in search.

Introducing integration depth

Switching from model B1 to D1 in Table 12 brings us to the third step of our analysis; from investigating AI use, to the breadth of exploration, and now the depth of eventual AI integration. We measure this integration depth as the share of key departmental processes that are supported by live AI tools – as the result of a principal component factor analysis, normalized to values between 0 and 100. A 1-point coefficient in the regression results can

[11] As described above, the age of the organization is used as the Heckman selection variable and is no longer appears as a correlate in the second stage of the model

thus be interpreted as a 1-percentage-point increase in integration depth. Comparing model D1 to the B1 reveals how the drivers for exploration and integration are similar and how they differ. A first difference is that previous organizational success now shows significance to a 5%-level. While slack resources therefore appear to have no effect on exploration breadth, they do facilitate a stronger integration of AI tools into live processes.

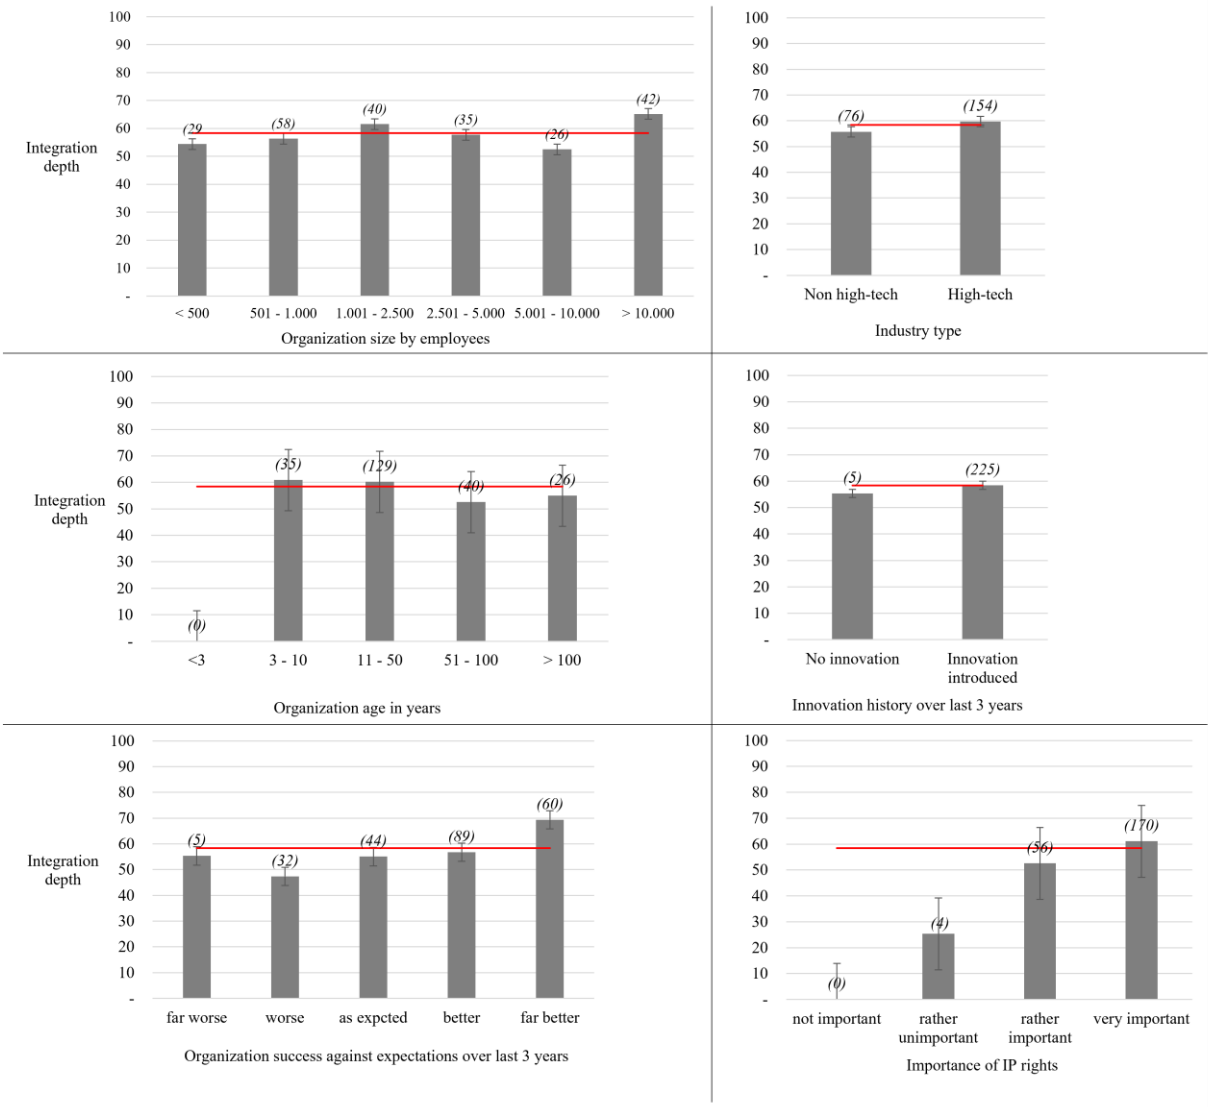

Figure 9: Effects of selected drivers on integration depth – for organizations using AI (number of observations per category in brackets above bars, total n=230)

Not immediately clear from the positive regression coefficient in Table 12, but revealed through graphic representation in Figure 9 is the interesting U-shaped effect that increasing success appears to have. Organizations with far better performance are associated with integration levels far above average. Organizations with worse performance show the lowest depth, far below average. However, organizations with performance *far* worse than expected have integration depth on par with much better performing organizations. While the low number of observations in the lowest category again means that we cannot draw clear causal inferences from this data, the connection is reminiscent of the concept of problemistic search (Cyert & March, 1963; Simon, 1957), where performance below aspiration levels triggers a search for solutions – in this case, perhaps in the form of more strongly supporting the introduction of new technology into organizational processes

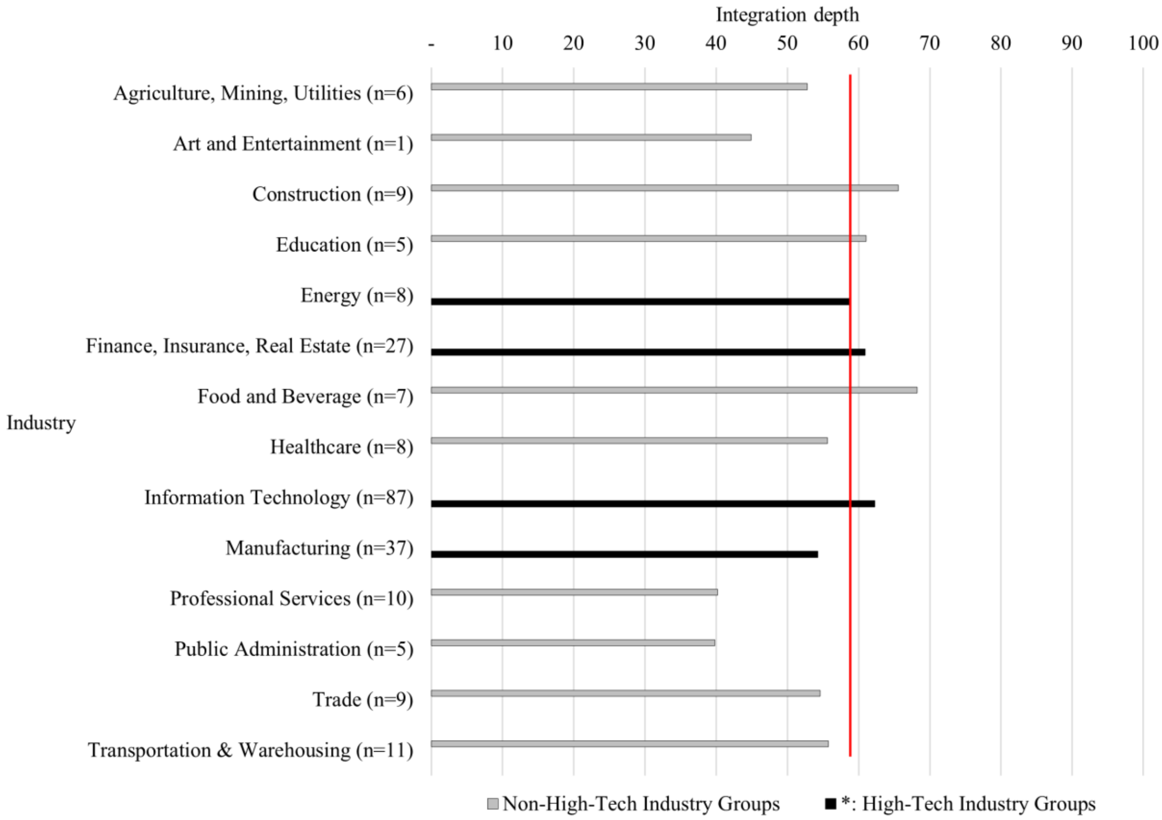

Figure 10: Average integration depth by industry clusters, indicating high-tech industry groups

A second interesting difference between exploration breadth and integration depth is that high-tech association is now no longer significant. We can explain this (lack of) result intuitively. Being part of such an industry may create either institutional pressures to engage with new technologies (Zucker, 1987) or allow well-trained employees to begin individual experiments. This makes organizations in these fields more likely to start working with AI, as the regression on AI use shows. Should organizations from a non-high-tech industry chose to begin working with AI against the odds, however, these organizations already differentiate themselves from their industry-peers, operating and committing resources as if they were a high-tech organization. The actual industry background no longer matters for integration success. Figure 10 shows the result visually, with the average integration depth for organizations using AI across various industries, classified by high- or low-tech, where high-tech fields are not more visibly engaged.

Perhaps more surprising than this are the similarities between exploration breadth and integration depth – which distinguish them from AI use. Innovation performance and the importance of IP rights both affect AI use but yield no significant results on AI integration. Differently from industry affiliation or size, these innovation variables reflect lived organizational behavior, rather than static and unchangeable characteristics – which we would have intuitively believed to affect the willingness to integrate new technology. Figure 9 shows these results and reminds us of the value of visual data representations on top of regression models. In the graphic, we see a clear relation between the treatment of IP rights and the depth of AI integration; organizations where IP rights are considered to be rather unimportant clearly lag behind other organizations in their ability to achieve deep AI integration. However, the number of observed organizations in this category is extremely low with only 4 responses – likely resulting in a loss of explanatory power in the regression model. While visually enlightening, the observation has to be weighed carefully. Of course, absence of evidence is not evidence of absence, but regression non-result in these innovation-focused factors that were the focus of many previous studies, are surprising.

Identifying further correlates

In the first step of our analysis, we have checked known drivers of AI use in our dataset. In the second step, we have tested whether these same drivers also act as significant correlates for the breadth of experimentation in organizations that do use AI – and then whether these also subsequently drive integration depth as organizations begin implementing AI into live processes. We now move beyond this set of previously analyzed drivers to add several other structural factors to the regression models, shown in the lower part of Table 12.

For exploration breadth, we indeed find two more significant regressors that add marginal exploratory power to model B1. First, the share of the share of departmental budget spent on digitization measures is significant and shows how an extra percent of digitization spend leads 0.005 more use cases. Second, stronger decentralization, in the form of increasing the hierarchical distance to the board of management by one level of hierarchy, shows an increase in the number of pursued use cases by 0.07. Based on existing literature, one could assume that this relationship is moderated by the size of the organization; for larger and multi-unit organizations, decentralization away from the headquarters or board of management could lead to reduced administration oversight (Brynjolfsson, Jin, & McElheran, 2021) which could in turn increase flexibility or investment freedom. Smaller organizations, on the other hand, could require more central attention to free up the required resources. The interaction term between decentralization and organization size, however, remains insignificant in our model B2.

Looking at model D1, for integration depth, a number of other variables turn out to be significant, showing clear differences in success drivers for exploration and integration. This forms a further important contribution of this study: the weight of explanatory power with regard to integration depth comes from departmental-, rather than organizational structures. For example, a single percentage increase in the departmental focus on creative tasks and the departmental investments in digitization lead to a roughly 0.28- and 0.24-point increase in AI integration depth in model D1. Similarly, the availability of data in digital formats in the department has a strong and highly significant effect, with an increase of one category in data availability leading to a near 9-point increase in integration depth. This matches what we know about AI use and complementary assets, that “AI use further appears alongside other high-potential technologies such as cloud computing and robotics, as well as building on the presence of digitally stored information (i.e., data) within firms" (McElheran et al., 2024, p. 397) – but adds further specification to these findings in that the required organizational invest must translate down to the departmental level.

Further down the table, we see another interesting finding: the use of organizational support structures. Such structures, measured as a binary indicator, lead to a 17-point increase in AI integration depth – the largest single effect of any measure. The use of strictly external structures, such as governmental programs or external consultancies reveals no significant results, which leads us to believe that the significance emerges from the introduction of comparable *internal* structures. Attributing significant organizational resources to the development of structures such as internal IT programs dedicated to AI, digitization incubators, or exchange network, can return valuable assets in the organizational quest for effective AI use.

Finally, the variables relating to structures of decision-making yield no significant results for integration depth, either. With both the organization size and the decentralization distance showing no significance, neither does their interaction term in model D2. The interaction turns significant only in model D10, where we additionally introduce the time of organizational AI use as a squared term – a complex four-way interaction to be investigated further in the next section.

Model corroboration

We use a number of corroborating models to check for variations in these results (see Table 10 for an overview). First, we focus on the measure of decision decentralization in models I. Instead of measuring the hierarchical distance between the board of management and the level of decision-making, we introduce the measure of board involvement – a binary variable to indicate whether or not the board of management was involved the identification and control of AI use cases (see Table 22 in the appendix for details on model I). The results from this model are very similar to those of the original model D, both in terms of significance levels and coefficient sizes.

The results of the next corroboration model R tell a similar story (see Table 23 in the appendix for details). Here, we create a new variable called "Department AI readiness", through a principal component analysis on the three originally significant variables all broadly related to complementary assets and AI readiness; the departmental focus on creative tasks, the investment into digitization, and the availability of digital data. Unsurprisingly, this variable remains highly significant, with a coefficient of 0.55 on a normalized scale of 1-100 against the existing dependent variable of integration depth. Slightly more surprisingly, model R shows the importance of IP rights being significant to a 5%-level for integration depth. A one-category increase in the importance of IP rights is now associated with a roughly 5% increase in integration depth – supporting our initial assumption of a positive correlation between the overall integration-focus of an organization and its intensity of AI use that was missing from models B and D [12].

Structural effects of time and decentralization

As the last element of this study, we present an initial exploratory analysis on the relation between organization size, the time of AI use, and the decentralization of decision-making, with the hope of inspiring future discussion. Returning Table 12, we see a statistically significant effect in the four-way interaction in model D10. While this complex four-way interaction is not easy to interpret, it does indicate the presence of our expected findings, based on previous research around organizational search and the creation of mental representations around new technologies (Boumgarden, Nickerson, & Zenger, 2012; Huber,

[12] Tables 24 and 25 in the appendix, corresponding to models P and Y change the dependent variable of the model. Instead of the combined PCF variable of AI integration depth, these models focus on each of the individual constituent dependent variables; the coverage of departmental key processes with AI tools in model P – and the intensity of AI use as measured by the regularity of use in model Y. The results here change the coefficients more than previous corroborations, but do not change the overall impression of the model.

2023; Kretschmer & Khashabi, 2020; Siggelkow & Levinthal, 2003). To further understand this interaction term, we attempt a margins analysis to plot average marginal effect curves[13]. Figure 11 shows the resulting graphic, plotting the average marginal effect of an increase in one level of decentralization on integration depth – across the previously defined categories of time. To include the suspected effect of the organizational sizes, we have grouped the datapoints into organizations with less than 1,000 employees (the blue/ bottom curve – typically representing small and medium-sized enterprises) and organizations with more than 1,000 employees (the red/ top curve), as well as showing the overall effect (green/ middle curve).

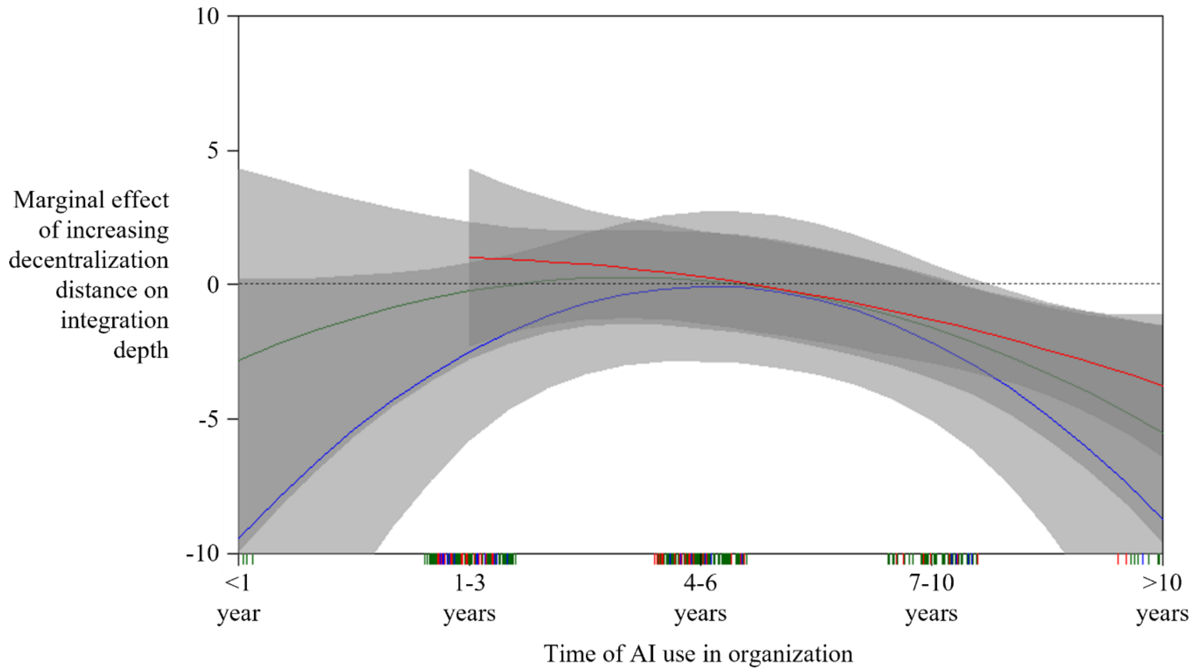

Figure 11: Average marginal effect of increasing decentralization in AI-related decision-making on AI integration depth over time.
Decentralization measured as levels of hierarchical distance between AI decision-makers and management board. Margins analysis using two-stage Heckman regression with robust standard errors. *Legend: Red – Large organizations of over 1000 employees. Blue – Small organizations of fewer than 1000 employees. Green – all organizations*

[13] Using the R library *margins* along with the *cplot* command, based on only the outcome regression of the two-stage Heckman model, with an adjusted covariance matrix for robust standard errors.

The graphic shows that, while the overall effect in the regression model may have been significant, neither of the two size-groups produce significant outcomes on their own. Nonetheless, as a first indicative result, we can observe some clear differences in the shape of the curves. For smaller organizations, decentralizing early in the AI journey negatively affects the AI integration – while for larger organizations, the opposite may be true; early decentralization appears to have a neutral or even slightly positive effect. The curves meet in the time-category 3, representing 4 to 6 years of AI use. Here, additional decentralization appears to have a neutral effect on integration. For even more experienced organizations, with more than 6 years of AI use, additional decentralization once again appears to have a negative effect, with a more pronounced effect for smaller organizations – still insignificant due to very large standard errors.

To further corroborate this indication, we perform the same margins analysis on model I10 – where the linear decentralization measure was replaced by the binary measure of board involvement, resulting in Figure 12. The logical connotation of this graphic supports the initial analysis: for SMEs of under 1,000 employees (blue/ U-shape curve), we had previously found that stronger *decentralization* had *negative* average marginal effects in the early stages of the AI journey, then turning neutral in time category 3, before turning negative again with more experience. In this version of the graphic, we find that board level involvement, that is, stronger *centralization*, has a strongly *positive* effect for SMEs – which turns marginally negative at time category 3, before turning positive again. For large organizations (red/ inverted U-shaped curve), the suggested effect is more pronounced and more clearly distinguishable than in the previous analysis: stronger board involvement, representing stronger *centralization*, appears to have a negative average marginal effect in the early stages, which turns positive around time category 4 (with 7 to 10 years of experience), before turning negative again.

This matches our initial intuition: Large organizations require free exploration, followed by strong intervention to create systems and standard operating procedures, but ultimately decentralized responsibility. Smaller organizations, on the other hand, require early input and motivation from top management; they otherwise suffer from limited resources, lower levels of organizational slack, and stronger pressures from daily tasks. Of course, even with these surface-level results matching those from previous investigations and the intuition based on literature, the broad shaded areas of standard errors across all curves once again show that these results are to be carefully scrutinized.

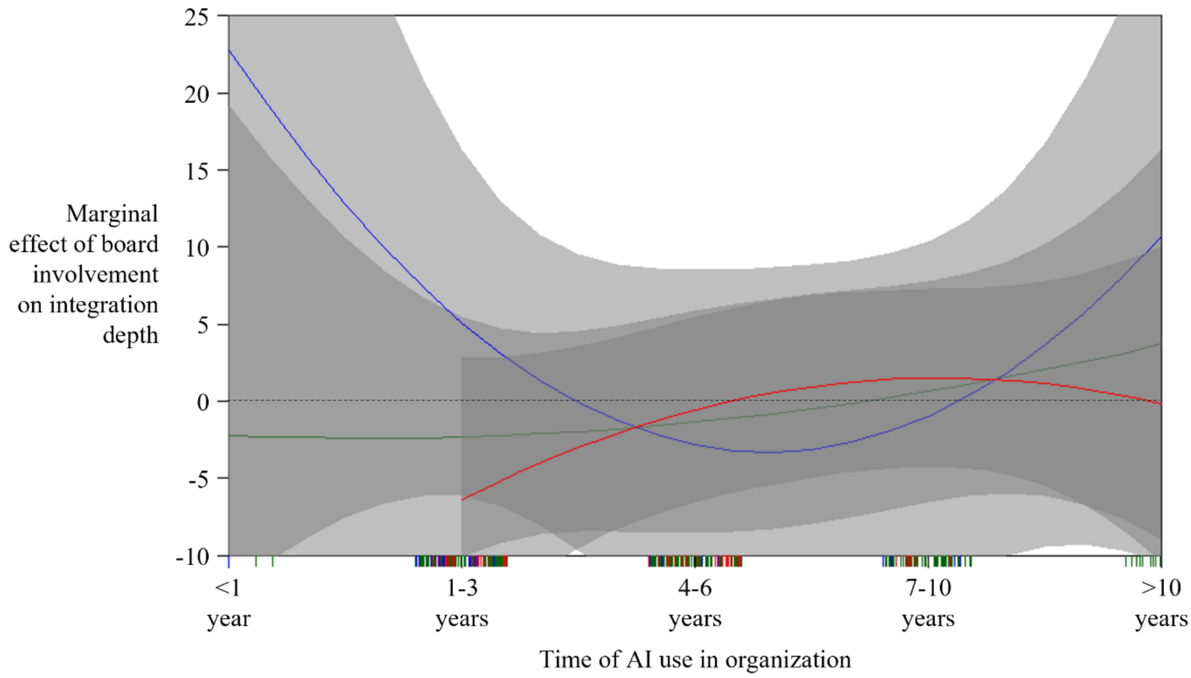

Figure 12: Average effect of board-level involvement on AI integration depth over time. Board involvement measured as binary influence on AI-related decision-making. Margins analysis using two-stage Heckman regression with robust standard errors. *Legend: Red – Large organizations of over 1000 employees. Blue – Small organizations of fewer than 1000 employees. Green – all organizations*

DISCUSSION AND CONCLUSION

By introducing to the conversation two novel measures of AI integration, exploration breadth and integration depth, we strengthen the collective understanding of how firm characteristics and structural choices affect the use of AI in organizations. Conducting a purpose-build survey allows us to engage more deeply with the sample organizations and reveal insights on a micro-level that were previously only looked at through a more macro-lens, and challenge established wisdom on AI integration. Our contribution therefore lies in asking the more specific question of “who uses AI *well*” – rather than “who uses AI” or “what is AI used for”. This opens a new stage of empirical research around AI in organizations.

Using descriptive- and regression analyses on binary AI use, as have previous studies, we find that use is directly related to organizational performance, history of innovation, and focus on IP rights (Czarnitzki, Fernández, & Rammer, 2023; McElheran et al., 2024; Rammer, Fernández, & Czarnitzki, 2022). We further find that AI-using organizations are more likely to be from high-tech industries and younger than non-using organizations when controlling for size.

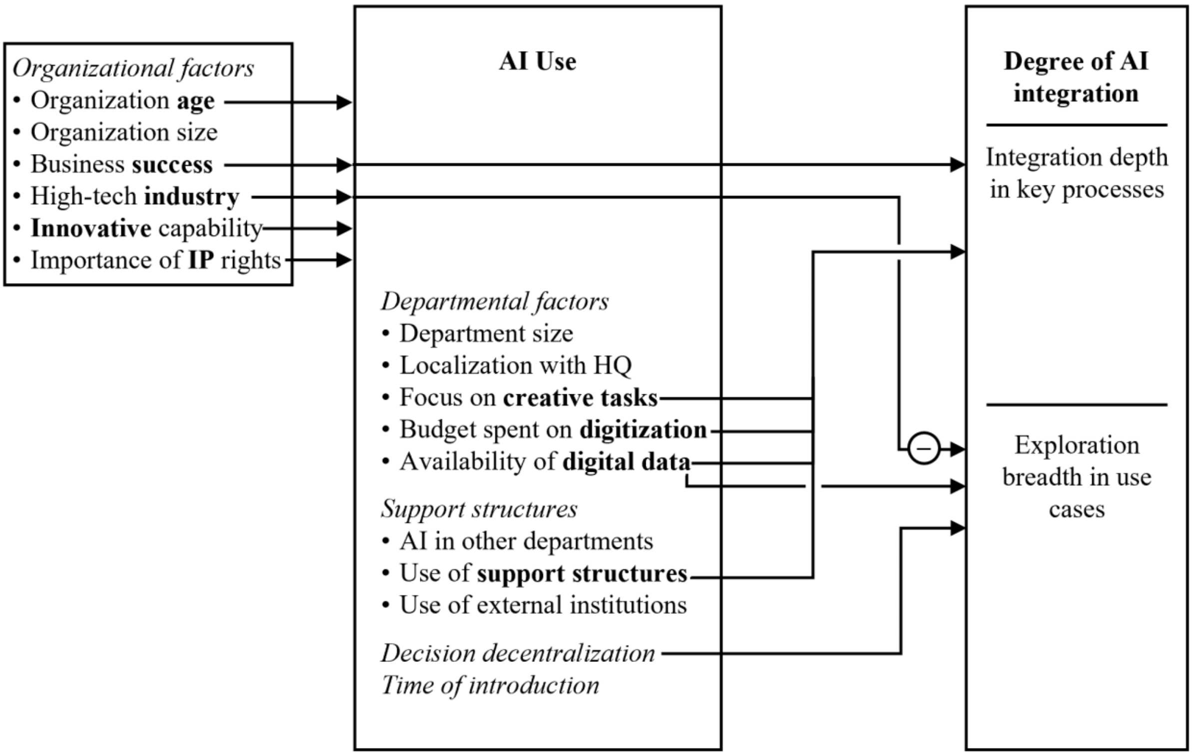

Figure 13: Results of Heckman regression analysis; significant effects of organizational and departmental drivers on the degree of AI integration by depth or breadth

Surprisingly, we find that of the above five drivers of AI *use*, only the association to high-tech industries affects AI exploration breadth, while only previous organizational success reliably affects AI integration depth. We summarize these effects visually in Figure 13 and further find that investments in digitization and decentralized decision-making both significantly contribute to exploration breadth. Yet, we also find that breadth appears to be strongly determined by factors outside the scope of our analysis; with an adjusted R-square between 10% and 16% for all model specifications, other factors contributing to our overall understanding remain to be identified.

For integration depth, we find that significant explanatory power rests on department-level factors, rather than the organizational-level factors that were previously often the focus of research. In line with such prior studies on the organizational level (Brynjolfsson, Jin, & McElheran, 2021; Brynjolfsson, Rock, & Syverson, 2019), departmental investments into complementary assets through digitization measures and the availability of data in digital formats significantly enhance AI integration. We specify this finding by revealing that these factors only drive depth, not breadth, and add to the discussion two more relevant factors: the departmental focus on creative tasks and the availability of support structures.

To summarize, our study reveals significant differences between the drivers of binary AI use on an organizational level, the drivers of AI exploration breadth as one measure of usage intensity, and the drivers of eventual integration depth into key business processes – a distinction that serves to inform future research and demonstrates the importance of using specific terminology in this discussion.

We further show how fundamental and unchangeable organizational characteristics play an important role in AI integration, but how other drivers are heavily reliant on top management and local leadership via various avenues of resource allocation. Specifically, we point out how departmental- rather than organizational-level factors may play a key role in supporting AI integration depth, focusing on creativity of core tasks, the availability of complementary assets through digitization and data availability investments, and the implementation of internal supporting structures.

Strategic implications

The distinction between fundamental and likely unchangeable organizational characteristics on one hand, and more malleable aspects of structure on the other is important to consider when deriving managerial or strategic implications for organizations. An organizations' size, age, success, or industry association, for example, cannot be easily changed, but can lead to behaviors which are more imitable and affect the organizational treatment of new technology.

The effect of business success, for example, could be interpreted as providing organizational slack, that is, available free resources that foster greater experimentation (Nohria & Gulati, 1996). While success is itself not easily replicable, organizations regularly undergo planning and prioritization processes. Prioritizing the allocation of resources to AI and other digitization programs as a strategic long-term initiative may be necessary, even in the face of potential constraints. To imitate slack effectively, however, it is important that these resources not be directly allotted to tools or expenses, but also allow free experimentation amongst interested employees (George, 2005; Mount et al., 2024; Vanacker, Collewaert, & Zahra, 2017). A similar approach may address three other key variables in our model: the budget spent on digitization, the availability of digital data, and the availability of support structures. While the two latter factors are not merely a budgeting exercise, but also require targeted restructuring, hiring, process-redesign, or procurement of advisory services, the intentional financial prioritization of AI-supporting structures promises to significantly accelerate adoption.

Two other variables present a more challenging puzzle to managers; the notions that younger organizations are more likely to adopt AI and that departments with a larger share of creative work are more successful at deep AI integration. Here, we challenge managers to identify specific features that these drivers may entail and attempt to model behavior after them. Younger organizations do not adopt AI just because they are younger, but more likely because they exhibit other traits that facilitate the adoption of AI. These could be a focus on talent acquisition, higher willingness to adjust existing processes, or more entrepreneurial risk-taking – which older and more established can chose to (carefully) imitate. The focus on creative tasks in departments deeper AI integration is perhaps even more surprising. While we know that modern large language models can improve creative performance (Dell'Acqua et al., 2023; Zhou & Lee, 2023), we also know that AI tools performing best in well-structured environments, which we would expect in more process-oriented departments. Especially before the advent of large language models, making AI immediately applicable to all language-based creative work, departments focusing on core processes such as supply chain or sales strategy would have been first in line to benefit from AI-driven analysis. Rather than in the tasks themselves, the reasons for the better depth of integration in departments with a creative focus might then lie in the culture and behavior of those departments. For example, a higher willingness to experiment with new tools and finding new approaches to problem-solving. Attempting to spread such approaches across organizations might therefore increase overall integration depth and eventual success.

Further research avenues

The analysis presented here opens several opportunities for expansion and further research. Firstly, the comparatively small sample size and method applied in the survey-based data collection bring with it a series of limitations. Especially for the non-AI-using organizations, we had to exclude a set of variables from the final analysis that could yield promising results in a larger-scale replication. These include the budget for digitization or the focus on creative- vs. process-tasks, which would make for interesting comparisons to the AI-using organizations in the final sample. Amending the analysis with a larger sample size of both users and non-users would therefore not only strengthen the reliability of the findings but also provide more opportunities for theorizing the department-level relation between drivers of AI use and integration depth.

Secondly, we do not differentiate between different types of AI solutions or applications areas. While AI is still a novel set of technologies and approached by many organizations as a strategically relevant monolith, more experience will lead to a more differentiated treatment of AI solutions. Identifying types of algorithms, such as supervised versus unsupervised learning tools, or machine vision versus language models, and investigating differences in their respective treatment along the dimensions presented here could yield further relevant insights. A related aspect is the missing consideration of the goal of AI introduction into the organization; using AI for product- or service-improvements can imply different requirements from using AI as a tool or process-technology (Huber & Alexy, 2024; Kiron & Schrage, 2019). Further differentiation here with a larger *n* could show how different structures may be beneficial or required for each type of goal.

Finally, we believe that the analysis of changing organizational structures to facilitate AI integration over time promises to develop into a highly interesting area of research – based on established ideas in theory around technology introduction in general and AI in particular (e.g., Boumgarden, Nickerson, & Zenger, 2012; Iansiti & Nadella, 2022; Siggelkow & Levinthal, 2003). From this literature, we know that the structural journey organizations undergo as they begin to grapple with new technologies is not linear but may instead move through various iterations of structures and power allocations through hierarchies over time.

We call for further research in this area, for example in the form of a panel dataset or other measures, to diving into these promising avenues, strengthening the link between existing knowledge on technology introduction and the continued importance of effectively managing AI.

Chapter 4:
Breakthrough Innovation and the Asymptotic Rationality of Artificial Intelligence

---

Note:
Previous versions of this chapter were presented at the 2022 Strategic Management Conference (Huber, Reetz, & Alexy, 2022a), the 2022 Vienna Conference on Strategy, Organizational Design and Innovation (Huber, Reetz, & Alexy, 2022b), and the 2021 Soph.I.A. conference (Huber, Reetz, & Alexy, 2021). This chapter is co-authored by Oliver Alexy and David K. Reetz.

"GPT […] wrote a thoughtful answer that was probably better than most of us in the room would have given. The whole experience was stunning. I knew I had just seen the most important advance in technology since the graphical user interface."
- Gates (2023)

"ChatGPT and its brethren are constitutionally unable to balance creativity with constraint. They either overgenerate […] or undergenerate […]. Given the amorality, faux science and linguistic incompetence of these systems, we can only laugh or cry at their popularity."
- Chomsky, Roberts and Watumull (2023)

Scholars have presented two primary explanations of how organizations achieve breakthrough innovation, that is, how they may develop novel technologies toward new-to-the-world products, processes, or services (Ansari, Garud, & Kumaraswamy, 2016; Bremner & Eisenhardt, 2021; Eggers, 2014; Hannah & Eisenhardt, 2018; Hargadon & Douglas, 2001; Zuzul & Tripsas, 2020). The first explanation emphasizes a frequentist logic: given humans' inability to predict a breakthrough, organizations are well advised to maximize variance by conducting as many attempts as possible to find one. For example, firms may enable decentralized innovative activities among employees (Gambardella, Khashabi, & Panico, 2020; Garud, Gehman, & Kumaraswamy, 2011; Rosenkopf & Nerkar, 2001), hoping that one of these attempts will turn out to be successful. In contrast, the second explanation emphasizes the importance of human ingenuity: breakthroughs are assumed to begin with some humans – mavericks, visionaries, or quirkies (Chai, 2017; Felin & Zenger, 2017;

Schilling, 2018) – somehow seeing things differently. Central to this second explanation is that, even though breakthroughs are still expected to result from recombination of prior knowledge (Gavetti, Levinthal, & Rivkin, 2005; Katila & Ahuja, 2002; Levinthal, 1997; Nelson & Winter, 1982), specific humans are thought to hold idiosyncratic knowledge to help them determine what may become a breakthrough, implying they may somehow foresee appropriate selection criteria.

With artificial intelligence (AI) entering the context of firms searching for breakthrough innovation (e.g., Amabile, 2020; Haefner et al., 2021; Hartmann & Henkel, 2020; Verganti, Vendraminelli, & Iansiti, 2020), the potential and limitations of this emerging technology has also been described through these two lenses. On one hand, in line with frequentist logic, AI can raise the scope and speed of variance generation in knowledge work drastically, even rendering some tasks in the creative process redundant (Dell'Acqua et al., 2023; Raisch & Fomina, 2021; Raisch & Krakowski, 2020). With an abundance of large language models and image-creative tools emerging over the past months and years, many observers (as Bill Gates above) praise the generative skills of AI and suggest that its abilities are nothing short of a revolution in creative work and potentially even suitable to identify (i.e., select) breakthroughs. On the other hand, in line with the ingenuity perspective, others argue that AI-enabled breakthroughs still "requires human understanding of the situation" (Agrawal, Gans, & Goldfarb, 2019, p. 91, Pearl and Mackenzie 2018) as well as "sensemaking" (Verganti, Vendraminelli, & Iansiti, 2020, p. 212), so that high-impact decisions, like selecting what may be(come) a breakthrough innovation, may never be fully delegated to algorithms (Camuffo, Gambardella, & Pignataro, 2023; Felin & Holweg, 2024).

These conflicting projections raise an interesting question around the role of humans and AI in the future development of breakthrough innovation: if both the possible scope of application and computational speed of AI continue to rise, what allocation of roles between human and algorithm will result in the near future?[14] We suggest that this question has two parts: first, to what degree may AI become able to emulate, if not copy or even surpass, innovation-related skills widely believed to be innately and uniquely human? And second, even if AI were not able to replicate these skills, (when) might we expect the remaining human-AI difference to matter?

To tackle these questions, as others studying AI and innovation before us, we draw on the behavioral theory of the firm and its metaphor of innovation as search (Cyert & March, 1963), in which boundedly rational actors seek performance improvements (see also Csaszar & Steinberger, 2021). We define a breakthrough as radical improvement in how customer needs are met.[15] Both the frequency and ingenuity perspectives agree that breakthroughs are rarely found in task environments known to firms, but rather as firms traverse into new ones, unearthed by novel technologies (i.e., new solutions) or previously latent customer needs (i.e., new problems). This process, we note, may be facilitated through practices allowing actors to access new task environments systematically (e.g., design thinking) and to convince others that an environment and the solution it contains is appropriate (e.g., marketing). Yet, if such techniques can be taught to humans, they may also be teachable to machines, which, as those machines improve, may steadily decrease the share of activities within breakthrough innovation for which humans hold a reasonable advantage over AI.

We capture our argument in the notion of *asymptotic rationality*: the more, and the more reliable information AI can acquire, the larger its role (relative to human actors) in developing breakthrough innovation. The more AI approaches rationality – which seems attainable in *known* task environments – the more it may not only contribute by generating variation, but by taking over the supposedly human task of selection. In turn, for as long as

[14] For the purpose of this analysis, we are focusing on AI systems in the near future – that is, even after the recent meteoric rise of large language models in 2022 and 2023, systems before the possible emergence of a general AI (Bostrom, 2016; Tegmark, 2017)
[15] This definition encompasses existing ("needs are met better") and new markets ("previously unmet needs").

humans continue to hold (but not express) idiosyncratic preferences and knowledge – as likely in *unknown* task environments – a role for human agents in selection will remain. Yet, by emulating human techniques of engineering breakthroughs, we suggest even this human advantage in selection may become increasingly smaller: essentially, AI may learn the same techniques that humans apply to access unknown task environments. First, AI may learn to imitate approaches to metacognition, such as using analogies, to produce novel and internally consistent ideas. Second, AI may be trained to conduct micro-experiments to learn to navigate even in previously unknown task environments. And finally, AI may simply be trained to redefine unknown environments into known ones – in essence, to develop the perfect marketing campaign to force-fit known problems and solutions.

We derive three contributions to literature around search and organizational innovation. First, we address current research emphasizing the importance of human intuition (Agrawal, Gans, & Goldfarb, 2019; Camuffo, Gambardella, & Pignataro, 2023; Gigerenzer, 2023; Verganti, Vendraminelli, & Iansiti, 2020) and highlight how the potential of AI to displace much human activity in the quest for breakthrough innovation may already be larger than acknowledged. In knowable task environments, AI may already match most humans' ability to create and select breakthroughs through both local and distant search. In unknown task environments, AI may emulate human thought processes to approach or at the very least support breakthrough innovation. The distinction between human and AI innovative capability then hinges on how much of human preferences remain unarticulated.

Second, we re-emphasize an important dilemma for practice: the dangers of over-reliance on AI tools reducing organization-specific capabilities (Beane, 2018; Candelon et al., 2023; Choudhury, Starr, & Agarwal, 2020; Zhou & Lee, 2023). With AI taking over routine and 'easy' tasks in the innovation process, new human experts may lack the training they require for judgement and reflections on the ethical use of technology, creating considerable social risk (Bengio et al., 2024). We emphasize that firms need to develop theories of value

(Felin & Holweg, 2024; Felin & Zenger, 2017) to guide their exploration of unknown tasks environments – which will also guide their use and deployment of AI for breakthrough innovation.

Finally, we highlight how the increasing prowess of AI tools point to a need not just for crisper theorizing, but also for cleaner terminology and more clearly stated assumptions

HUMAN CREATIVITY AND INGENUITY IN BREAKTHROUGH INNOVATION

Before we study how AI may impact the quest for breakthroughs, we shed light on current perspectives of how human agents arrive at them, and extend that by an AI view thereafter.

Bounded Rationality, Innovative Search, and Breakthrough Innovation

Rationality, following Simon, is "concerned with the selection of preferred behavior alternatives in terms of some system of values whereby the consequences of behavior can be evaluated" (1997, p. 84). March, in turn, added that rational choices are "derived from model-based anticipation of consequences evaluated by prior preferences" (2006, p. 202). This also implies that "a rational choice process requires an a priori act of problem framing and representation before the execution of a rational choice calculus can be carried out." (Levinthal, 2011, p. 1517). In empirical reality, agents only show bounded rationality: they do not have access to and cannot comprehend or process all relevant input factors simultaneously so that decision-makers often lack the "complete knowledge and anticipation of the consequences that will follow on each choice" (Simon, 1997, p. 93).

Building on Simon, the Carnegie School sees innovation as process and outcome of searching for solutions to pre-identified problems (Argote & Greve, 2007; Cyert & March, 1963; Gavetti et al., 2012; Knudsen & Srikanth, 2014; Nickerson & Zenger, 2004; Simon, 1997). Bounded rationality, in turn, poses a natural barrier to the kind and degree of innovation human agents may identify: while human actors may continuously improve their understanding of the real world by searching, it may remain impossible that they ever fully comprehend it. Accordingly, they are thought to satisfice, that is, settle for solutions to their problems that are good enough rather than optimal (Simon, 1955).

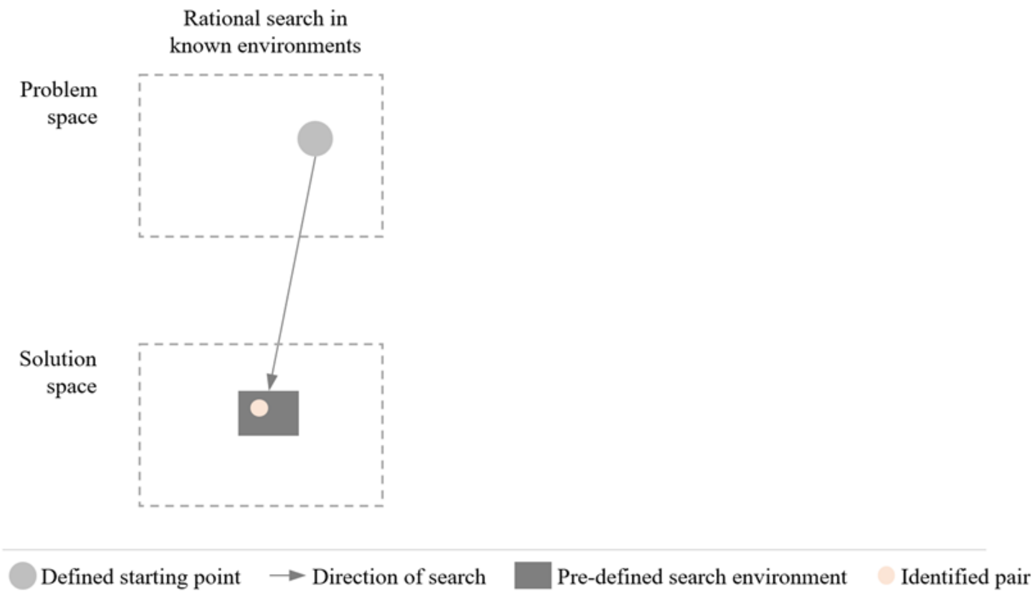

Figure 14: Process of search with clear starting point and pre-defined solution space

For as long as searching agents are boundedly rational, the efficacy of search is driven by two factors – the blades of Simon's metaphorical scissors: the cognitive abilities of the searching agent, and their fit to the structure of the task environment, that is, the environment in which search is conducted (Simon, 1990). Current work rarely problematizes the task environment (see, e.g., Posen, Leiblein, & Chen, 2018 for a detailed analysis), a fact that may surprise anyone who is asked to count the number of quiet people in a room (cognitive abilities) with their eyes opened or closed (the task environment). Rather, the task environment is usually assumed as benign, knowable (or even as known) so that boundedly rational actors can acquire a reasonable understanding of the problem space, which encompasses all aspects of the problem and how those are perceived by potential stakeholders, such as customers (Figure 14). Given such an assumption, even boundedly rational actors may express well-specified problems, from which they can derive clear-cut goals to direct search (MacAulay, Steen, & Kastelle, 2020; Simon, 1996). For example, in the often-used metaphor of NK models (Levinthal, 1997), actors are assumed to be able to define a problem in form of a payoff landscape, and then assess the performance produced by the efforts they exhibit toward solving it: in essence, agents can measure their ‘altitude’ on a coordinate system given the inputs they chose.

In turn, whether such search leads to a breakthrough is an innate feature of the landscape, determined by (known) technological limitations and market preferences. If these are stacked favorably, agents may be able to find breakthroughs simply through local search, by continuously identifying minor improvements over their current performance. In the NK metaphor, bounded rationality implies limited visibility from the point at which the agent starts these efforts. Whether this new location is a local or a global optimum would depend on where the agents started their search and how far they can see: even if a landscape may contain a breakthrough innovation, actors on a linear development trajectory may miss the right recombination of existing factors to produce a breakthrough if it is too far away from their current focus – a pattern similarly observed in empirical reality (Chai, 2017; Fleming, 2001). Breakthrough innovation through such local, linear improvements may occur when individuals manage to find a particularly high-performance peak on the landscape of their current task environment. Kneeland, Schilling and Aharonson (2020) provide the example of BioSteel and super-resolved fluorescent microscopy as illustrations of how continuous incremental steps may accumulate to breakthrough innovation.

To overcome potential limitations of local search, such as getting stuck on local optima, individuals can explore the performance landscape through distant search (Levinthal, 1997). In essence, by purposively picking one or multiple additional starting points for search far away from where the agent currently is located (through, for example, an exploratory experiment) and then searching locally from those, agents may identify various peaks on the landscape, and select the most favorable one. In that context, analogies may often play a role in allowing agents to enrich their mental models of an existing task environment and identify new starting points for subsequent local search efforts (Gavetti, Levinthal, & Rivkin, 2005).

Merrill Lynch`s ‘financial supermarket’ business model has been described as a breakthrough innovation resulting from such distant search (Gavetti & Menon, 2016).

Moving to New Task Environments to Create Breakthroughs

While the importance of search in known task environments is beyond doubt, a significant body of work argues for alternate sources of breakthroughs. Breakthroughs often require doing or seeing things differently, meaning that agents somehow have a different perspective on technology or market needs. In the language of search, these agents are essentially working in a *different* task environment, in which they may find breakthroughs no one else can see (see Figure 15) (Adner & Levinthal, 2008; Gavetti, 2012; Gavetti & Menon, 2016).

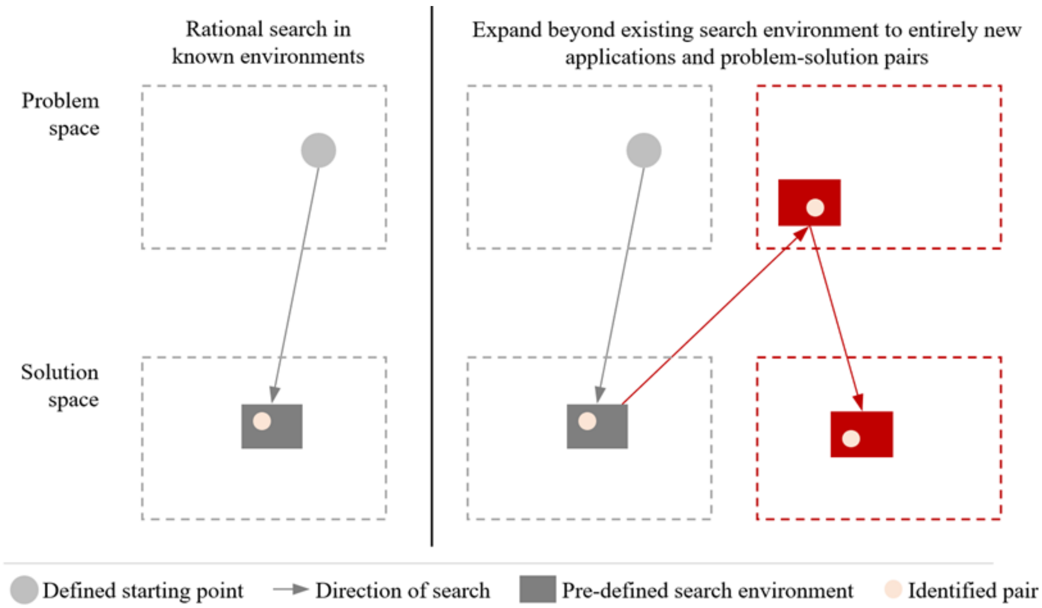

Figure 15: Capabilities of Human Intuition in breakthrough innovation outside of known environments

The logic of user innovation (von Hippel, 1986) provides a particular poignant example. Users engage in innovation because they foresee a particular *idiosyncratic* benefit from its outcome. While others may also benefit from the outcome, user innovators go on to produce it because their need is particularly pressing, and they have the skill and resources to get started. Successful examples of this type of user-driven breakthrough innovation are the heart-lung machine, developed over 17 years by John H. Gibbons, or the dishwasher, developed by Josephine Cochrane, frustrated at her porcelain breaking from the stress of manual scrubbing.

Exaptation (Andriani, Ali, & Mastrogiorgio, 2017; Andriani & Cattani, 2016) follows a similar logic. Here, existing knowledge or technology is redeployed for a different use –by accident, as in the example of the development of the microwave oven, or on purpose as in the case of Velcro modeled after burrs. Exaptation, too, enables searching agents to identify novel market needs to be addressed by novel technological solutions, and hence leave the existing task environment in the hope to find a breakthrough.

As these examples also indicate, stories of breakthrough innovation following an ingenuity approach usually feature individuals envisioning some kind of future performance landscape, along with the breakthrough innovation and bottlenecks therein, and launching experimenting efforts to overcome those.

To this end, humans are thought to have two kinds of capabilities which allow at least some of them to insert some purposive direction when they search for breakthroughs.

First, given their own preferences, experience, or talent, some individuals may be able to re-imagine the recombinatory potential of resources. As in our above examples, such users may interpret the generative affordances a technology grants them to address a largely idiosyncratic problem (Bailey et al., 2022). Second, for these ideas to become breakthroughs for reasons other than sheer luck (Denrell, Fang, & Liu, 2014), individuals need foresight (Gavetti, 2012; Gavetti & Menon, 2016): a sufficiently accurate prediction of the behavior of others (in particular: customers) faced with unknown phenomena in an unknown context.[16]

[16] An example of such latent (i.e., existing, but as-of-yet unobservable and unknowable) market needs are the hypothetical future customers who cannot express their needs for a product, as captured in the adage ascribed to Steve Jobs: “It is not the customers job to know what they want.” Such undefined task environments must lack clear boundaries by definition Simon (1957). They contain latent unarticulated preferences (March, 1978), beliefs (Felin & Zenger, 2009, 2017) and values (Rindova & Martins, 2017) that give rise to new courses of action, all of which is infinitely difficult to codify in advance.

Yet, as Winter (2012) highlights, just like luck itself, such idiosyncratic factors that equip agents with an innate better starting points cannot be the basis of theories trying to explain effective *organizational* search processes, especially those hoping to generate breakthroughs (somewhat reliably). The only recommendation that would follow is that firms should try to hire the smartest individuals, which – even if we assume it would be possible to spot which individually-held vision is right in advance – seems a trivial point to make at best.

Structured Organizational Search for Breakthroughs in New Task Environments

We suggest that infusing the ingenuity approach to breakthroughs with a frequentist perspective may return some agency to searching firms, and matters particular to study the role and importance of AI: While genius and luck may be hard to come by or tell apart (Denrell, Fang, & Liu, 2019), firms may design or draw upon processes and structures through which searching agents try to identify new task environments systematically. For example, the literature suggests that firms may attempt to find a mix of a frequentist and ingenuity approach by deploying several, possibly even competing projects across strategic focus areas, each formed around ideas by individuals or small groups with an internally consistent vision that is externally aligned with the firms' own. In this context, 3M, Google, and others have long deployed deliberate processes to maximize their odds at, essentially, enforcing serendipitous encounters by creating circumstances through which human agents, individually or collectively, recombine problems and solutions from various domains to set the stage for potential breakthroughs around the firms' key technologies (e.g., Bock, 2015; von Hippel, Thomke, & Sonnack, 1999).

The literature describing how to design organizations that foster serendipitous encounters that can lead to breakthrough innovation describes three complementary design elements: first, there are (infra)structural factors linked to the physical layout and department structure of the firm (Allen, 1970; Tushman & O'Reilly III, 1996). Second, motivational factors such as granting autonomy are known to foster more creative work (e.g., Amabile, 1996). Finally, procedural factors imply firms deploying internally consistent sets of search activities through which human agents may approach breakthrough innovation more systematically, of which the literature has identified three broad sets: (i) meta-cognition, (ii) experimentation, and (iii) framing. Figure 16 depicts these practices schematically.

Meta-cognition implies that actors apply a supposedly inherently human, learnable skill of identifying parallels or similarities between task environments (e.g., Dahl & Moreau, 2002). Analogies, associative mental models, heuristics, or simple rules are primary examples of this approach (Davis, Eisenhardt, & Bingham, 2009; Eisenhardt & Bingham, 2017; Gavetti, 2012; Gavetti & Menon, 2016; Gigerenzer, 2023). Through such meta-cognitive techniques, individuals transpose a known problem structure onto an unknown problem, and then treat it as if they were in a known environment. The quality of such an approach, of course, depends on the quality of the analogy, which is only knowable post-hoc. For example, while imagining becoming either the Google or Apple of interstellar travel are perfectly valid analogies based on which agents could design a search effort to realize a breakthrough innovation, there is no reason to believe one analogy is per se better than the other. Still, drawing on such a meta-cognitive process may assist search actors in structuring their efforts to become internally consistent (Siggelkow & Rivkin, 2009; Sull & Eisenhardt, 2015).

Micro-experiments share many similarities with meta-cognition – in particular the goal of learning about the true nature of the novel task environment over time. Yet, the focus is less on extrapolating from agents' existing knowledge to the new context, but on quickly identifying the key features of the new context: agents iterate between potential definitions of a problem (e.g., unexpressed user needs) and potential components of a solution (e.g., novel technological affordances) to identify what combination of problem and solution could create a breakthrough. To do so, agents may initiate search both with and without prior theories about which problems are important or which solutions are valuable.

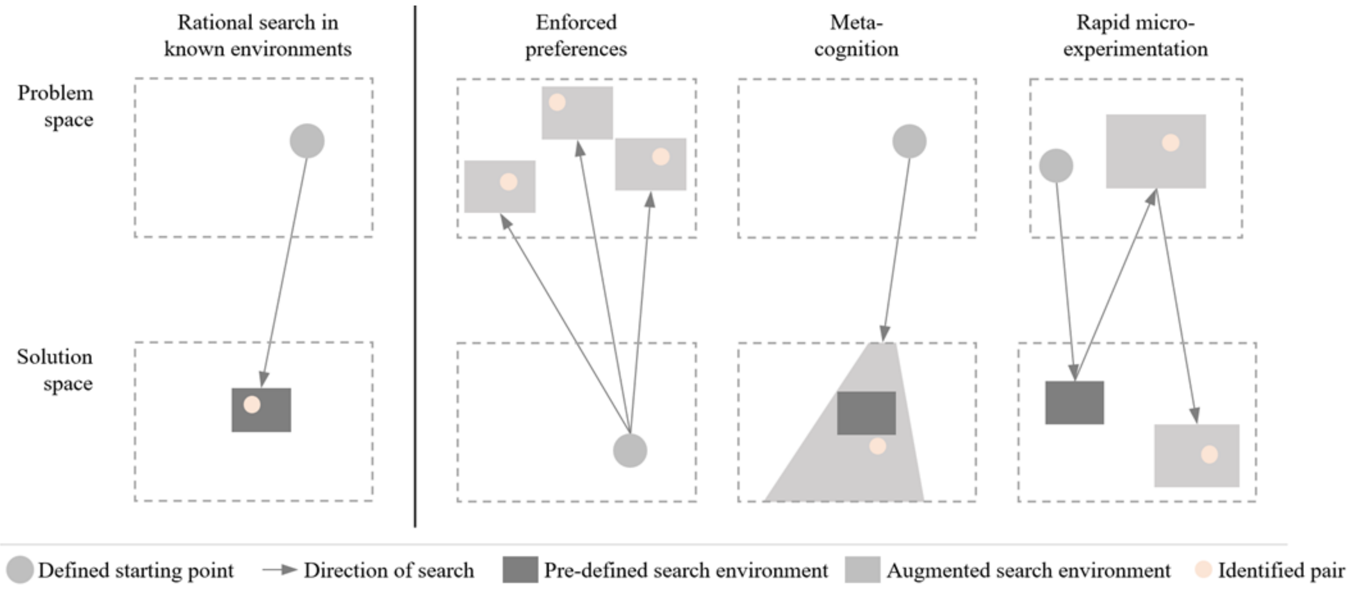

Figure 16: Human Cognitive processes enabling breakthrough innovation

Following either principles of the scientific method or pragmatism, they may experimentally develop and refine theories about the nature of their environment abductively or inductively (Camuffo et al., 2020; Felin et al., 2024; Felin & Zenger, 2017; Zellweger & Zenger, 2023). Examples include popular techniques like design thinking or the lean startup. Such structured approaches allow agents searching for breakthroughs to deploy processes that may render their efforts less error-prone; in particular, these methods often aim at minimizing the costs resulting from overcommitting to likely failing courses of action (i.e., avoiding errors of commission) (e.g., Adner & Levinthal, 2004). Yet, they still cannot guarantee that any of the chosen options will turn into a breakthrough.

Finally, rather than finding the true nature of the task environment and developing a breakthrough solution for it, firms may try to *frame* the task environment so that solutions they can or will be able to produce are seen as more favorable. Beyond marketing, such efforts include various market (e.g., pricing) and non-market strategies (e.g., lobbying, partnering) to create conditions that increase the odds of a product being declared a breakthrough. For example, the winners of races for eventual technological standards often were not the better technologies, but those who were able to exploit latent (i.e., previously unexpressed) market needs. As such, efforts at shaping a task environment to make a known product or solution become a breakthrough through improved perception by consumers is a highly promising strategy – its eventual effectiveness, however, still remains questionable if multiple actors try deploying such efforts at the same time (Gavetti, Helfat, & Marengo, 2017; Helfat, 2021).

In sum, breakthrough innovation may result from various agentic search processes. In line with traditional descriptions of search, breakthrough innovation can be found through linear optimization by newly re-combining known factors to identify previously unknown solutions within an existing task environment. In turn, what is described as distant search seems to confuse two distinct concepts; a distant solution to a clear-cut problem in a known task environment, versus reimagining the problem itself to create a different or altogether new task environment in which to look for breakthroughs. In this latter approach, even non-genius agents may draw on structured processes in the hope of being guided toward developing novel and internally consistent solutions. The external fit and, hence, breakthrough potential of these solutions is once again determined by the (unknown) structure of the new task environment, which can only be identified post-hoc. Alternatively, human agents may try to shape the task environment itself, by influencing existing or latent preferences of (future) consumers.

THE ROLE OF AI IN ORGANIZATIONAL SEARCH FOR BREAKTHROUGHS

When John McCarthy and colleagues coined the term artificial intelligence in their proposal for the 1956 Dartmouth Conference (McCorduck, 2004), they hoped "to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves" and that "the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving" (McCarthy et al., 1955, pp. 2, 13). From those first ideas of machines as thinking or acting humanly, advances in machine learning have led to AI being discussed as thinking and acting *rationally* (Norvik & Russell, 2021). That said, AI, too, can only act rationally under limited information: it depends on available data. In early expert-systems, such data would need to be "elicited from human domain experts and painstakingly hand-coded in a formal language" (Bostrom, 2016, p. 8). Through dramatic improvements in computing power, costs of data storage, and data availability through the internet, machine learning systems have become viable, and increasingly capable of analyzing and adjusting to data without the need for human intervention and re-configuration – though they still depend on data that is available (Brynjolfsson & McAfee, 2014).

The key difference between humans and AI is hence often argued to lie in the ability to generate *novel* data. Even for non-geniuses, the human advantage has often been described as having intuition, the ability to make connections and generalizations based on a limited number of observations (Felin & Holweg, 2024; Gigerenzer, 2023). Given their ability to process multiple types of information – including emotions of others – in a context-aware fashion, humans, accordingly, will outperform machines when it comes to breakthroughs. Still, beyond the idiosyncratic advantages conveyed by luck or genius, the process of human imagination operates remarkably similar to algorithmic search: using previous experience as training data, humans match potential solution ideas to their mental models of newly imagined pay-off structures, making "the output of a data analysis run by the human brain […] feel like inspiration" (Ludwig & Mullainathan, 2023, p. 1).

Accordingly, approaching the level of ingenuity of *most* humans may largely be reduced to a challenge in data availability and processing capacity: how the machine is trained to learn from data, and the degree of context awareness it may have.[17]

With respect to training, we may distinguish three kinds of how AI learns from data: supervised learning (where algorithms require pre-labeled datasets for training in order to predict specific values), unsupervised learning (where patterns are detected in unlabeled

[17] See earlier chapters for a more detailed analysis of these concepts and their interactions with each other.

datasets to create clusters), and reinforcement learning (where a model optimized a specific reward function in uncertain environments based on direct action feedback). Regarding, an algorithm's ability to recognize its environment, we may distinguish four types (Hintze, 2016): reactive machines (as the baseline for systems incapable of holding information, used for single-shot tasks), limited memory AI (addressing most systems in use today, where selective data is saved and made recurrently available to use for future decisions), systems possessing theory of mind (able to understand the state of mind of other agents), and fully self-aware AI (recognizing itself as an agent capable of thought and emotion). While self-aware AI can safely be attributed to the domain of science-fiction for now (Bostrom, 2016; Tegmark, 2017) current-generation AI, large language models built on neural networks and reinforcement learning, are slowly being thought of as possessing theory of mind (Kosinski, 2023; Schossau & Hintze, 2023; Xu et al., 2024).

In turn, whether and how AI can produce innovation will depend on the quality of the data it has available that matches its algorithmic design and the degree of understanding it has about its environment, where an AI that is less advanced on one dimension may even overcompensate on the other. For example, assuming an AI existed in the 1860s, had been fed with all knowledge existing of that time, and tasked to identify new models of transportation, it would not have required particular context-awareness or updating-capability to discover the automobile. Yet, the less information it would have had, the more it would have needed context awareness and updating capabilities not to merely identify faster horses. More formally, even simple supervised learning tools such as decision trees could achieve variety by following clearly defined logics of association, while more complex reinforcement learning tools could build on feedback provided against a clearly defined reward function to create solutions increasingly different from the known environment.

In sum, as a baseline, we expect both non-genius human agents as well as AI agents tasked to create breakthroughs to face similar constraints around the availability of sufficient and sufficiently accurate data. What will determine the relative performance of humans and AI will hence most likely not be their ability to ideate based on the data that is available (i.e., how fare relatively in known task environments), but if and how they can produce or obtain access to new, reliable data based on which breakthroughs may potentially be developed (i.e., how they fare relatively in unknown task environments).

AI as Facilitator of Breakthroughs in Known, Well-Specifiable Task Environments

Indeed, an almost trivial consequence of our baseline argument is that AI will be faster than human agents at identifying breakthroughs within a knowable set of solutions to a specific problem in a known, well-describable task environment. Therein, human ingenuity will matter less compared to the frequency at which AI can create and evaluate potential solutions.

The application of AI in drug discovery provides a fitting example. For decades, humans outperformed computers at identifying the most energy efficient three-dimensional resting structure of known proteins, a piece of information vital for drug development, making use of tools like Foldit – a gamified software sustaining a world-wide community of hobby-protein-folders (Khatib et al., 2011). Now, AI neural networks with reinforcement learning have learned to be more performative as this task (Eisenstein, 2021). Similarly, AI is seen as a vehicle to deliver finally on the promise of high throughput screening, by which a (very!) large library of compounds is tested against a drug target to identify a starting point for a potential drug (Chan et al., 2019).

We also note how AI need not necessarily be taught the structure of a knowable problem in advance – as long as there is sufficient information specifying the problem, unsupervised or reinforcement learning AI can find latent rule-like structures itself. Not only do clustering algorithms for fraud-detection work in this manner, but consider also how AI has a stellar record at board and video games, such as Chess, Go, or StarCraft.[18] In such

[18] Computers are reliably beating humans at chess at least since the Cray Blitz working with logics- and tree-based algorithms gained Master status in 1981 (Hyatt, 1981), at Go, with a exponentially larger set of options

systems, the algorithm can learn over time to recombine a set of known options into high-quality solutions that eventually match the underlying rules. Following a similar logic, advanced large language models such as ChatGPT and others can recombine existing ideas, including existing mistakes, to create what resembles creative thought in essay form, computer code, or even recipes (Knight, 2022). Similarly, image creation algorithms such as Dall-E or Midjourney use existing image libraries and previous user requests to build and re-use a database of image components, creating novel solutions (images) to any problem (prompt) composed in a known or similarly knowable search space.

In sum, for well-specifiable, knowable problems, AI is capable of quickly covering a larger array of the task environment much more quickly than humans. As such, AI has been empirically shown to out-perform humans in specific tasks of knowledge-work, such as the generation of a broad range of ideas (Dell'Acqua et al., 2023), or even the creation of artwork in both amount as well as creativity and quality (Zhou & Lee, 2023). Still, we note how it is average rather than ingenious humans who benefit most from working with AI: the performance of ingenious individuals sees little improvement, while the rest of the population begins to move closer to them (Dell'Acqua et al., 2023; Gaessler & Piezunka, 2023; Hartmann, Exner, & Domdey, 2023), in particular when deploying multiple AI systems jointly (Doshi et al., 2024; Webb, Holyoak, & Lu, 2023).

We therefore expect that, if a breakthrough innovation is still waiting to be found in a well-specifiable task environment, AI will find it at least at the same level of quality as human agents – AI might not necessarily be better, but it will certainly be faster. AI must therefore already be seen at least a fierce competitor to human agents for innovation challenges in which data on technological options and projections of future market needs exists (such as

and moves, since DeepMind’s much-discussed victory through a deep-learning neural network in 2016 (Chouard, 2016), and finally at the vastly complex StarCraft since 2019 (Garisto, 2019), where the rules allow for an even wider range of movements and strategies in a virtual fog of uncertainty.

historic customer preferences for products or designs, buying behavior in global markets, or expected competitor actions and reactions). In such settings, human ingenuity or creativity could be replaced by a learning algorithm, carried forward by continuous advancements in computing power and ever-better access to data. Ever-increasing codification allows for, in Simon's words, expressing innovation problems on a *"continuum of degrees of definiteness"* (1997, p. 183): Even though a large range of factors may be relevant, as each factor is sufficiently precisely defined, problem complexity decreases. The higher the degree of definiteness, the more useful artificial computational power may be to augment or even replace human capabilities (Newell, Shaw, & Simon, 1957; Newell & Simon, 1972). As such, where task environments can be defined or described in reasonable detail, AI may be able to combine variation and selection to identify a breakthrough as the true optimal solution.

Thus, if and only if the task environment is well-specifiable based on existing data, AI may be seen as capable of (almost) rationally identifying breakthrough innovation. Even if this data may originally be acquired through managerial expertise and stored informally as know-how, it will often be codified or codifiable to a sufficiently precise extent to allow for algorithmic analysis. Even ignoring a possible future existence of general AI (Bostrom, 2016; Tegmark, 2017), if a search problem can be well-specified (meaning AI can find 'the' problem and predict reliably not just what should be done, but how others will respond to its choices), AI systems may already be understood as overcoming constraints of human rationality, in so far as we understand the system to "select alternatives which are conducive to the achievement of the previously selected goals" (Simon, 1997, p. 4).[19]

---

[19] Of course, this kind of AI-powered search is only rational within the confines of its own world, defined by the scope of (a) the data available and (b) the predefined problem, which together make up the representation of reality in which the algorithm operates. The rationality of the AI may therefore be considered limited by the mental models of the actors defining the problem which the AI is employed to support.

AI-guided Breakthroughs Outside of Known Search Environments

As argued above, the strength of human ingenuity is commonly seen in our ability to identify promising ideas outside of known search environments. Supposedly different from machines, humans may for example draw on intuition or tacit knowledge to imagine ways of addressing unstable and unknown preferences (March, 1978): predicting what future preferences *could* look like and which ones might be (more) relevant to successful search is central to human cognition (Peirce, 1957). Drawing on imagined preferences like predicted future markets addressed via solutions from envisioned technological development, some humans may create breakthroughs by designing new-to-the-world problem-solution pairs. AI, on the other hand, is usually thought incapable of identifying idiosyncratic and latent needs, which would imply that breakthroughs in new search environments may only originate from human thought.

Theoretically, of course, this answer is likely correct: someone, somewhere may always hold idiosyncratic knowledge allowing for new-to-the-world recombination with the potential for a breakthrough. Until true generative AI, that would imply that humans will most likely always surpass machines in their potential to create breakthroughs in new task environments. Accordingly, we believe that a much more important question is asking whether this difference in potential actually matters and will continue to matter.

As said above, from an organizational perspective, a world in which only geniuses produce breakthroughs will only lead to a war for those few individuals highly talented at anticipating the future. Beyond that, organizations will need to deploy motivational, structural, and procedural factors so that a larger number of employees may discover and develop new technologies and markets and try to merge those into breakthrough innovations.

How does AI come into this picture? From an AI perspective, a genius may well remain unpredictable. But where humans deploy well-specifiable mechanisms to improve their odds at achieving breakthroughs, AI might become a significant factor. Admittedly, motivational factors might not matter to a machine that can run 24/7 with no feelings or knowledge of self.[20] Structural factors include *how and where* to embed AI, and how to organize the collaboration between humans and AI, for example in ensembles (Choudhary et al., 2023; Jia et al., 2023; Puranam, 2020). Hence, aside from an algorithm that would beat humans at deciding which tasks to allocate to whom (work on this is progressing with regard to human and robotic actors, e.g. Chakraa et al., 2023; Zahedi, Sengupta, & Kambhampati, 2024), selecting good structural factors might improve organizational performance overall, but not tilt the human-AI balance. Of course, the AI in the human-AI ensemble may take over tasks previously done by human agents. For example, creative employees could draw on AI to generate variance, and possible help with pre-selection to keep a large set of ideas manageable (Jia et al., 2023).

However, structural factors alone cannot determine whether an AI may successfully take over those creative tasks that render humans supposedly (more) effective at breakthrough innovation: identifying what would be a valuable idea for a breakthrough (i.e., one worth selecting) in an unknown task environment.

While often ascribed to intuition, we have laid out how this part of the creative process can be facilitated through meta-cognition, micro-experimenting, and framing. These practices by which (non-genius) humans try to enter new task environments *systematically* are codifiable precisely because of their systematic, rule-like nature. In turn, if these practices may be learned and applied by an AI, we may also expect that AI may increasingly encroach on humans trying to identify breakthroughs even in new search environments.

Looking at *meta-cognition*, AI is well suited to generate new solution data by applying generative models, such as heuristics or means of abstraction, to new contexts (Schilling, 2017). By doing so, AI can easily search beyond pre-defined task environments and include previously unattractive solution alternatives as possible matches for a given

[20] Coincidentally, though, given the text corpuses based on which large language models are trained, phrasing them in a polite way has often been named a key recommendation for successful prompt engineering.

problem. This requires programmers to create AI systems which, rather than relying on sets of pre-determined algorithmic steps, allow the system to define the most suitable steps on its own. Rather than following patterns of movement, such systems follow heuristics to determine the patterns. This reflects what organization scholars have labelled simple rules (Davis, Eisenhardt, & Bingham, 2009; Eisenhardt & Bingham, 2017) or associative mental models (Gavetti, 2012; Gavetti & Menon, 2016) – once again making the AI emulate human cognitive processes. The latest generation of algorithms by DeepMind Technologies, for example, seems to be moving in just this direction; building on the game-changing alphaGo algorithm, the system MuZero is now being taught to maneuver in spaces where rules are messy and unclear a priori and have to be inferred by the system independently (Kelion, 2020; Schrittwieser et al., 2020a, 2020b).

Human agents also employ much less advanced creativity techniques to use existing data to identify breakthroughs via analogies. One example is TRIZ, the theory of inventive problem-solving developed by Genrich Altshuller (Gadd, 2011). TRIZ tries to break down engineering challenges into foundational problems, or contradictions. Once these problems have been identified, inferring possible approaches from past successful inventions, the tool leads to high-level solutions or principles. Clearly, such a data-heavy, highly algorithmic approach could easily be supported by AI: notably, Altshuller originally generated his principles by studying thousands of highly inventive and successful patents solving contradiction. In turn, an AI tool might identify more reliably than a human which contradiction is present in a focal situation (by drawing on more data), and automatically generate and pre-validate fitting solutions.

Clearly, such meta-cognitive approaches can be extended to allow for *micro-experimentation*, which we had defined as iterating quickly between potential definitions of a problem and potential components of a solution to identify key features of a new context. Indeed, this is precisely the process by which AI tools such as DeepMind have learned games such as Chess, Go, or StarCraft. In these games, AI generates new data by independently searching for problems and solutions, inferring one from the other, and using results as training data for the next iteration. simultaneous discovery of novel problem-solution pairs.

A prime application for AI in experimentation lies in training the AI to emulate human cognitive processes of design (Hatchuel, 2002; Le Masson et al., 2019) or imagination (Alvarez & Porac, 2020; Rindova & Martins, 2021) by replicating structured processes approaches in this area, such as the lean startup, design thinking, or speculative design. There, in line with our prior arguments, only two kinds of human inputs will definitely remain necessary for the foreseeable future for two reasons. First, humans would somehow need to express their preferences, based on which AI would engage in its design task, so that AI would roughly know where to search. In essence, AI needs to be given some form of direction or goal. Second, as long as data about problems and solutions is incomplete, humans may be required in exercising judgment about the output the AI produces (Allen & Choudhury, 2022; Choudhury, Starr, & Agarwal, 2020) – in essence, deciding whether an AI-generated suggestion is useful and, increasingly more important, ethical. For example, an AI may develop a prototype, receive human feedback, learn from that, and so on. All activities of alternative or variance generation may be supported or even conducted entirely by AI. More precisely, while the some kind of scope for the AI effort has to be given by an instructing agent, the AI may conduct any subsequent experiment autonomously – possibly even on a real market. And with the AI increasingly learning about market preferences, it may also learn to produce reliable predictions on what of the variance it has produce might most likely be considered a breakthrough by a market.

We see the biggest immediate area in which AI could take over human effects in *framing*. We suggested that, as for humans, the main challenge for AI in breakthrough innovation is the impossibility of reliably predicting unknown or unstable preferences. Yet, rather than guessing those, AI tools may be used to *define* them and push them onto specific groups of consumers – in plain language, AI could be trained to generate the perfect marketing campaign. In the vocabulary of innovative search, AI could first help identify which problems (i.e., known needs) could be addressed by a given solution (i.e., a technology, product, or service the firm has). In a second step, especially when possessing theory of mind and thereby leveraging the ability of understanding customers' mental states to appeal to emotions, AI would generate new demand for the already-existing solution through personalized and targeted advertisement campaigns. The story of Cambridge Analytica, analyzing and shaping citizens' political preferences using comparatively straightforward dimensionality reduction algorithms (Hindman, 2018), provides a case in point. As with experimentation, the role of humans would shift to selecting which problems or solutions will be communicated to consumers, but the design and execution of the campaign could largely be left to AI. However, as the Cambridge Analytica example also highlights, with increased reliance on AI, ethics also must move to the forefront.

In sum, our arguments around the role of AI in facilitating breakthroughs in unknown environments are similar to those in known ones: the more data AI has available to determine what is possible and valuable, the more it may be seen as capable of predicting breakthroughs. Does this automatically imply that machines now systematically and reliably produce breakthroughs by the hour? As of the time of writing, no: humans still have a natural advantage in producing reasonable approximations of what is plausible in new search environments for which insufficient data exists. However, as the human *techniques* to do so in a structured fashion may increasingly become accessible to AI, at least in theory, an ever-increasing share of the breakthrough-generation process may be taken over by machines. Importantly, these tasks go beyond those for which the application of AI is widely accepted – support in generating variance in the creative process. Even today, AI can help make important selection decisions, or at the very least help enforce them as a marketing support tool.

DISCUSSION

Reconciling Perspectives: AI, Breakthrough Innovation, and Asymptotic Rationality

We started our paper observing conflicting predictions around the effect of AI on the role of human agents in in breakthrough innovation, even though both the frequentist and ingenuity perspectives agree that AI may significantly impact this process in the future. Through a frequentist lens in particular, AI revolutionizes breakthrough innovation (see alsoRaisch & Fomina, 2024): given neither humans nor a machine could predict breakthroughs, the higher frequency at which AI may produce suggestions could render it a viable competitor. Through an ingenuity lens, if problems in known task environments can be sufficiently well-specified in advance, AI will also be able to take over a significant share of the creative endeavor. Finally, for unknown task environments, we note that non-genius humans rely on structured processes for innovation, which an AI may also be taught.

Accordingly, our arguments suggest that at this point already, AI may well be seen as more than a simple variance-generation machine. Rather, whenever sheer computational power improves information processing *and* allows for automated selection, AI may come close to identifying breakthrough in almost rational fashion. More broadly, when considering breakthrough innovation to be a recombination of known ideas, the only limit to AI's capabilities is the data it has been trained with and its computational capacity. Humans have exactly these same limitations – but in principle, only AI might be able to overcome them.

In contrast, when emphasizing the spontaneously ingenious component of innovation – presenting something new to the world as part of a breakthrough from previously unknown spaces – AI clearly can never reach the level of some exceptional humans. This is especially true given how humans may always hold unique, yet-unexpressed, and, hence non-codifiable information that only other humans may be able to access. Here, AI cannot reliably have sufficient information on what breakthrough it should produce for whom and hence cannot evaluate and rank the potential solution alternative it generates as well as such humans could.

Combining these ideas into a single framework, the role AI in breakthrough may best be seen as asymptotically rational: near-perfect in known environments with sufficient data, but inferior to ingenious human agents in breaking through to unknown ones (see Figure 17).

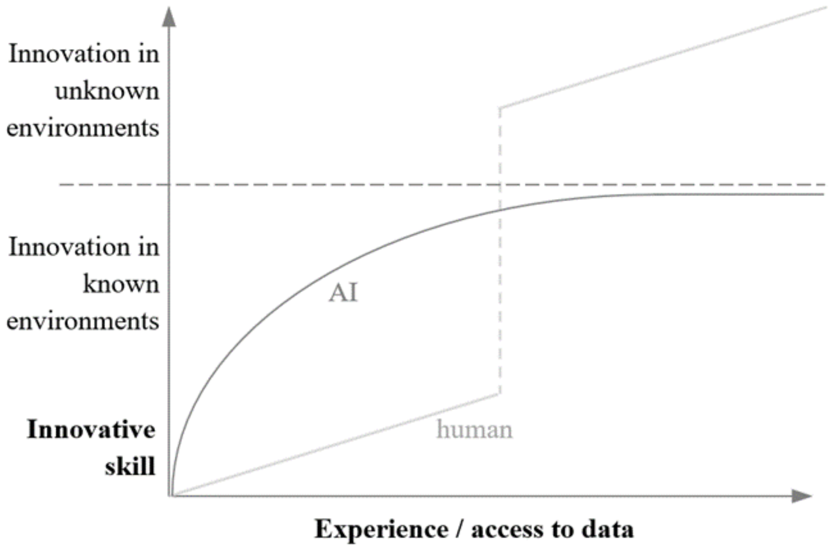

Figure 17: The notion of asymptotic rationality of AI

Our key point, however, is that there is a lot of important middle ground between those two extreme perspectives, and that the share of innovative work AI can contribute to the process of breakthrough innovation even in unknown spaces may already be larger than sometimes described. We fully subscribe to arguments suggesting that, until a potential artificial general intelligence emerges, some humans will always have some advantage over AI (e.g., Felin & Holweg, 2024; Gigerenzer, 2023). However, we also suggest that the *difference* between what humans and the AI can do as part of breakthrough innovation will steadily decrease. Indeed, the supposed human advantages around creativity, originality, intuition, or framing may hold out only for as long as we also ascribe these humans with exceptional, randomly distributed, and likely unpredictable qualities of being able to imagine non-existing task environments. AI could only perform on par with those if it was truly rational. For *most* human agents, on the other hand, the difference may become increasingly negligible. We will still require humans to articulate *where* to search, *which* solutions are ethically permissible, and, in unknown task environments, *which* out of a series of potentially viable options to select. The more data AI will hold even on those tasks, the smaller we should expect the advantage to be that even well-trained human agents may hold.

This is especially relevant if the focus of theories of search is on how organizations work. AI may already emulate the processes many human agents follow under such circumstances and take over a substantial share of work in the innovation process. For example, AI may take over most of the essential task of variance-generation of the creative process, with humans increasingly focusing on selection. In turn, the closer AI may come to emulating human creativity as a systematically codifiable recombinatory process, the more it will be able to develop selection capabilities – and therewith approach breakthrough innovation even in unknown search spaces, as shown in Figure 18.

In sum, rather than focusing on what AI may be able to do one day, we tried to highlight how it may already be able to contribute significantly to the development of breakthrough innovation. We do not suggest that AI will, at least for the foreseeable future, ever attain the full rationality to take over all of breakthrough innovation from human agents. Yet, as we try to capture in the notion of *asymptotic* rationality, as AI advances and more data becomes codified, the difference may increasingly matter less.

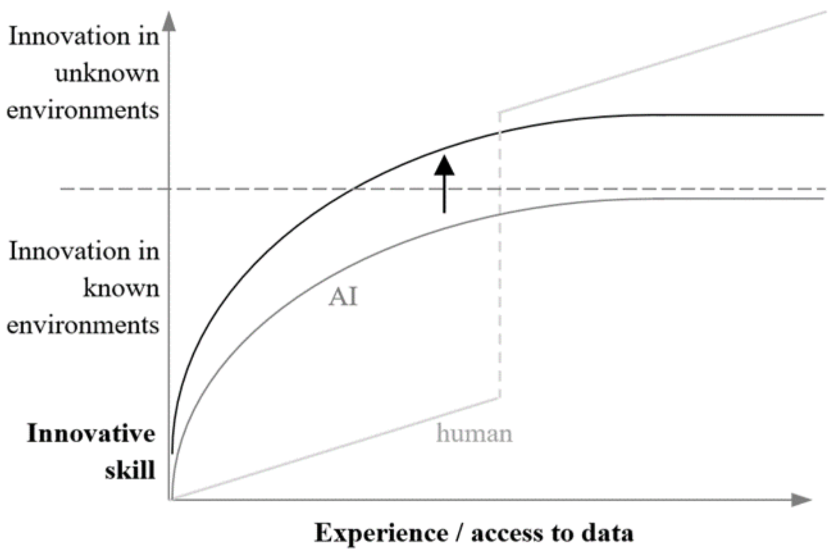

Figure 18: Asymptotic rationality of AI - Improving AI capabilities in innovation

Implications for Theory

Drawing on these insights, we propose three contributions to discussions of search and breakthrough innovation in literatures around organization, strategy, and innovation.

First, our insights speak to work discussing the organization of creative and innovative work, and breakthrough innovation in particular (Agrawal, Gans, & Goldfarb, 2019; Verganti, 2008; Verganti, Vendraminelli, & Iansiti, 2020). This work has particularly emphasized the human capability of selecting or framing problems, of merging insights from various domains, and, hence, of human intuition and capabilities of being quintessential to breakthrough innovation. Our discussion highlights that these arguments only apply in parts. A share of human ingenuity is untrainable – it is the result of luck or serendipity, and takes shape in the form of idiosyncratic, new-to-the-world resources, needs, or ideas (Busch, 2024). It is only this part that machine intelligence, for a considerable time, should be incompatible of reproducing, but it is not clear how humans might be able to reliably reproduce it, either (Schilling, 2018; Winter, 2012).

Rather, we suggest that a significant share of those human capabilities that make new task environments accessible and produce breakthroughs are learnable, trainable, and, hence, must also be codifiable. From the set of creative tasks, only judgment – selecting which problems are worth inquiring, and which solutions are worth following up – remains (Agrawal, Gans, & Goldfarb, 2019). Here, it is not clear what good judgment should be, a priori: regarding what will turn out to be a breakthrough, there may be procedurally sound judgment, but no prediction (Foss & Klein, 2012; Packard & Clark, 2020). As soon as procedurally sound judgment is inferable based on accessible data, the creative process can be executed by AI. Hence, human agents might be well-advised to focus their creative endeavors on selecting promising problems (Nickerson, Yen, & Mahoney, 2012) and submitting those to AI. In such a world, beyond identifying (ingenious) individuals with idiosyncratic world-view and testing those at large numbers with (frequentist) AI, it seems that the source of breakthrough innovation, and possibly competitive advantage altogether, could migrate to firms developing better theories about their environment (Felin & Zenger, 2017), and learning how to draw on AI to test those (Iansiti & Lakhani, 2020; Shrestha et al., 2021).

This argument, however, points to an important second insight: even if AI *could* replace a substantial amount of creative work in the innovation process, organization designers may still be well advised to use less AI support than is possible. If we assume the ability to select to be a learned capability resulting from having judged and possibly even generated 'simple' variance, then the over-eager outsourcing such tasks to AI will prevent human agents from reaching this level of mastery in the future.[21] Human agents may need such simpler tasks to be able to build up the capabilities necessary to understand what AI can and cannot do, exercise ethical judgment, and retain their ability to reign it in, especially in situations where the AI makes confident but potentially biased predictions. Importantly, these concerns also apply to setting in which AI is only supposed to be used as a variance-generating tool. Judging whether AI has produced what looks like an adequate depth or breadth of search is an evaluative task not dissimilar from evaluating if any of the problems or solutions identified are valuable.

Finally, our arguments contain interesting insights for discussions of search broadly. Drawing on the Carnegie School (Cyert & March, 1963; March & Simon, 1993), we laid out how most models of search assume local and distant search to occur within a given and knowable task environment, often modeled as an NK-landscape. In such a setting, however, AI does not only imply that searching agents can extend their scope of search (i.e., search at greater distance): The logic of asymptotic rationality we described implies that, in known

[21] This notion is supported by first empirical results, e.g. Allen and Choudhury (2022), who discuss the USPTO's decision to reduce AI use in prior art detection enable better examiner learning, (Candelon et al., 2023) showing that the diversity of ideas declined by 40% when knowledge workers used AI support, or (Zhou & Lee, 2023, p. 3) showing that the novelty of art declines when artists use AI, in an "expanding but ineffective creative space"

search environments, the distinction between local and distant search becomes obsolete as AI will be able to cover near-instantaneously the entire known search space. At the very least, this insight is a call for the use of more precise terminology in discussions of types of search and innovation, which features a proliferation of adjectives for clearly related phenomena (local, distant; radical, disruptive, etc.). Specifically, we note how these different terms, and varying uses in the literature result from different assumptions with respect to who searches for what, where, what is known, and what is knowable –that is, whether problems are assumed to be clear-cut, task environments are known, and problem and solution landscapes given, known, or knowable. As recent discussions around the topic of uncertainty have made clear, these thoughts go beyond nit-picking on theoretical jargon, but have foundational consequences for how we think about strategy and organizing (Packard & Clark, 2020; Rindova & Courtney, 2020; Rindova & Martins, 2021): only when a problem, based on a strong preference function, is defined can a solution be truly found – otherwise, as in the case of breakthroughs, problems, solutions, and their connecting preference functions are idiosyncratic (Adner & Levinthal, 2008) and need to be legitimized jointly (Grodal & O'Mahony, 2017).

Implications for Practice

For organizations, our results echo important insights around the need to enable members of organizations to familiarize themselves with AI. The recent hype around ChatGPT may have already led to a mindset shift in the general belief about the degree to which creative work may be impacted by AI at all. These developments clearly suggest that more mundane creative tasks – such as composing tweets or LinkedIn posts – could easily be taken over by an AI. Which tweets and posts to pick, however, and for what reason, is still a strategic choice of the firm. This strategic choice could become even more important as the availability of data in certain areas of organizing increases. Organizations might be tempted to follow the data – and focus their efforts on problem-solving or development in the use of AI in those areas. As a variation of the common streetlight effect (Freedman, 2010), this means that other, more pressing issues may be ignored in favor of data-rich fields. Critical observation of the implementation and use of AI therefore remains an important strategic task.

Policy needs to be aware of potentially unforeseen consequences that unregulated AI development and deployment may have (Bengio et al., 2024). Here, the arguments around the process we labeled preference shaping may be particularly relevant. This potential of AI to actively shape consumer preferences and markets has important ethical implications also in innovation and beyond (as our Cambridge Analytica example also highlighted), and also links to discussions of innovative dynamics and progress (Park, Leahey, & Funk, 2023). Indeed, corporations may reasonably shift their attention away from investing into uncertain R&D projects to identify new technologies, when they have better tools for commercializing whatever technology they have at hand. Policymakers need to be aware of such potential dynamics in order to be able to look out for what may otherwise turn out to be a source of market failure.

Limitations and Suggestions for Future Research

A core boundary condition of our argument lies in the availability and accuracy of data, and continuously advanced ways of processing it, especially as it relates to unarticulated preferences of population groups. Progress here seems to be faster than even the forecasts of optimistic analysts (Balasubramanian, Ye, & Xu, 2020). The diffusion of large language models may have rendered some of our arguments much less speculative than they may have seemed otherwise – displaying the previously unimaginable shifts of the asymptote we describe, albeit, for now, within known search spaces. Similarly, the first AI innovation support tools were presented, which begin to integrate the logics we present (Bell et al., 2023; Bouschery, Blazevic, & Piller, 2023; Just et al., 2023).

As prior work laying out how AI may improve the pace and output of innovation processes (e.g., Haefner et al., 2021; Raisch & Fomina, 2021; Seppälä, 2021; Verganti, Vendraminelli, & Iansiti, 2020), we encourage future work to look more into how AI may enable qualitative rather than quantitative change in the innovation process. For example, as initial work suggests (Kittur et al., 2019), given its sheer endless processing power, AI could be harnessed to connect entirely disconnected agents –even individuals rather than members of a traditional organization – to enable entirely new models of variance generation and idea selection in new types of innovation funnels.

Finally, future work may study what happens if the asymptote we describe is continuously pushed. In our view, answering this question will require a significant amount of fieldwork to sharpen our understand of how AI changes the nature and understanding of work (e.g., Beane, 2018; Beane & Orlikowski, 2015; Kellogg, Valentine, & Christin, 2020; Lebovitz, Levina, & Lifshitz-Assaf, 2021; Lebovitz, Lifshitz-Assaf, & Levina, 2022; Waardenburg, Huysman, & Sergeeva, 2022). Here, one goal would be to understand not just what organizations and creatives could do with AI, but to understand how and why specific practices might be adopted at a larger scale. Beyond fieldwork with organizations adopting AI, we see particular promise in working with organizations trying to develop and deploy algorithmic tools to be used by other organizations. In turn, shedding light on such organizations may not only give us a view at future technological affordances, but possibly also the strategy, ethics, and policy questions that come with those developments.

Concluding Remarks

We currently find ourselves in The Between Times of artificial intelligence – the time "between the demonstration of the technology's capability and the realization of its promise" (Agrawal, Gans, & Goldfarb, 2022b, pp. 3-4). These times are marked by great uncertainty that envelopes not just individuals, but also organizations. This uncertainty for organizations has become the subject of a growing field of academic research.

On a fundamental level, organizations must understand what AI means to them, which involves potential changes to their business- and operating models (Iansiti & Lakhani, 2020; Lanzolla et al., 2020), but also requires predictions on the highly uncertain outcomes of actions in the competitive environment (Krakowski, Luger, & Raisch, 2022; Townsend et al., 2024). On a process- and systems level, organizations may struggle to provide the required complementary assets (Berg, Raj, & Seamans, 2023; Brynjolfsson, Rock, & Syverson, 2019) and create a culture of acceptance and trust to enable effective integration (Glikson & Woolley, 2020; Hasija & Esper, 2022; Kawaguchi, 2020; Vanneste & Puranam, 2024). These factors become increasingly relevant as AI affects more and more areas of organizing, as previous research has shown; for example processes of coordination and control (Choi, Liu, & Shin, 2023; Giermindl et al., 2022; Kim, Wang, & Boon, 2021; Waardenburg, Huysman, & Sergeeva, 2022), daily routines of work (Beane, 2018; Dodgson et al., 2022; Lebovitz, Lifshitz-Assaf, & Levina, 2022; Pachidi et al., 2021), creative expression (Dell'Acqua et al., 2023; Jia et al., 2023; Zhou & Lee, 2023), or the constellations of teams (Choudhary et al., 2023; Murray, Rhymer, & Sirmon, 2021; Puranam, 2020).

In contributing to this wide-reaching conversation, I employ the perspective of The Behavioral Theory of the Firm, which views organizational change and innovation as a process of problemistic search by boundedly rational actors (Cyert & March, 1963; Simon, 1957). This allows me to address four key questions around organizations and the uncertainty they are exposed to through the emergence of AI, which I answer in the respective chapters of this dissertation:

1. What choices must strategic leaders make to prepare the organization for effective AI use?
2. How do organizations manage the early process of AI integration and how does it affect their structure and coordination of decision-making?
3. Which organizational characteristics and structural choices enable the broadest exploration and deepest integration of AI tools?
4. How does the increasing rationality of AI affect the process of organizational search altogether, as expressed through the roles of humans and AI in innovation?

In Chapter 1, I present a literature review that focuses specifically on the role of strategic leaders. I discuss how leaders face a fundamental choice between working with AI or towards AI – that is between using AI as a tool in organizational processes or as a core element of the organization's strategy and product offering. This is the elementary decision from which I derive further implications for organizational strategy, structure, and processes, as well as for the capabilities that are required of the strategic leaders themselves. In this analysis, I also point out why leaders and human agents in organizations will continue to play an important role in organizations, for two reasons.

Firstly, all AI systems, even modern large language models, are reliant on the quality of their training data (Agrawal, Gans, & Goldfarb, 2022c; Bogost, 2022; Chomsky, Roberts, & Watumull, 2023; Marr, 2023). While systems can be confident in portraying their results, they may be built on shaky foundations and can give wrong or even dangerous advice as the examples of Google's fantasizing AI assistant and Amazon's biased HR algorithm show (Dastin, 2018; Grant, 2024). Human domain expertise is essential, to remain aware of the limitations of automated prediction, capable of intervening at the right times, and sufficiently skilled to exercise judgement between alternative options. The importance of this factor is only multiplied in situations where ethical concerns are in play on top of those relating to business success (Fjeld et al., 2020; Torresen, 2018).

Secondly, humans may long remain the only actors capable of interpreting previously uncodified data and rare events. With AI tools relying on training data and available existing information, rare shocks are by definition not part of their immediate range of analyzed options – and so the recognition of existential dangers to business may still lie in the domain of human strategic managers for the foreseeable future (Agrawal, Gans, & Goldfarb, 2022c).

In Chapter 2, I analyze in more detail one of the aspects pointed out in Chapter 1; changes to organizational structure. The analysis follows recent calls pointing to a lack of insights particularly in the early stages of organizational decision-making (Joseph & Gaba, 2020), and highlights how, different from previous studies, the fundamental uncertainty surrounding AI requires organizations to create a shared mental representation of AI use while discovering the associated payoffs at the same time (Csaszar & Levinthal, 2016; Knudsen & Srikanth, 2014; Posen et al., 2018).

In this process, purposeful iterations of decentralized and centralized structures play a key role, representative of an ongoing re-allocation of decision-making authority. While previous literature has also pointed to the importance of dynamic structures in times of technological uncertainty (Boumgarden, Nickerson, & Zenger, 2012; Siggelkow & Levinthal, 2003; Smith & Lewis, 2011), this study presents a different perspective on the role of centralization: Where centralized structures are most often associated with the exploitation of new processes or technologies *after* the creation of shared mental representation, I show how centralization can also act as a necessary driver to allow for this creation in the first place – before exploitation takes place in more decentralized structures. While in many ways reminiscent of previous instances of technological change and fitting into the existing analyses of the Behavioral Theory, the introduction of AI brings with it unique challenges based on its technological characteristics – such as the required deep integration into specific domains. This requires structures more akin to platforms (Gregory et al., 2021; Kretschmer et al., 2022) with centrally established processes and tools, but decentralized decision-control, rather than the traditionally expected fully centralized exploitation.

Building on and directly complementing this newfound understanding of organizational structures aiding the introduction of AI is the quantitative analysis in Chapter 3. In this study, I show the necessity for a nuanced discussion in conceptualizing the use of AI in organizations. Previous work has largely built on either binary AI use as the dependent variable in analyses of organizational characteristics (McElheran et al., 2024), or as the independent variable to measure its effects on types performance (Berg, Raj, & Seamans, 2023; Czarnitzki, Fernández, & Rammer, 2023; Dell'Acqua et al., 2023; Rammer, Fernández, & Czarnitzki, 2022). I point out how this conceptualization misses an important intermediate step; namely, which organizations manage to integrate AI effectively. I begin by measuring AI use and corroborating existing findings, but then show how the organizational characteristics and structural choices affecting *use* differ significantly from the factors driving the *breadth* of AI exploration, and again differ from the factors driving the *depth* of AI integration.

Importantly, the organizational core characteristics that seem to facilitate AI use, such as the organization’s age or its association with a high-tech industry, are not the key factors driving AI integration. Instead, integration seems to be more related to factors malleable by managerial choices, such as the availability of supporting structures for departments employing AI tools, the share of departmental budgets reserved for digitization measures, or the availability of digital data.

Finally, Chapter 4 returns to the foundational principle of The Behavioral Theory, the assumption of bounded rationality, and explores the role of the apparent increase in rationality of AI in the process of organizational innovation. With this, I build on the metaphor of innovation as search (Cyert & March, 1963) and first analyze innovation in knowable and known search spaces. In these spaces, defined by a good understanding of the organizational environment, providing sufficient codified training data, AI may easily be able to outperform human decision-makers, who often lack the “complete knowledge and anticipation of the consequences that will follow on each choice” (Simon, 1997, p. 93). Indeed, in such an ideal setting, AI may come close to achieving rationality, defined as selecting “alternatives which are conducive to the achievement of the previously selected goals” (Simon, 1997, p. 4). The discussion, summarized in the framework of asymptotic rationality, becomes more complex as we turn to entirely new search spaces, where breakthroughs often take place, and where training data is scarce. In this scenario, unique skills are often ascribed to human actors, with the ability to understand latent and unarticulated preferences (Felin & Zenger, 2009; March, 1978; Rindova & Martins, 2021). In discussing these spaces, my co-authors and I concede that a select few humans may always be able to use their innate genius to create innovation (Gigerenzer, 2023; Verganti, Vendraminelli, & Iansiti, 2020) – but we argue that this is entirely unpredictable and so should not serve as a guideline for creating breakthroughs (Schilling, 2018; Winter, 2012). Instead, the more predictable and therefore more useful approaches to managing new search spaces for organization (namely meta-cognition, micro-experimentation, and framing) could already be adequately replicated using AI. With these capabilities only continuing to grow, organizations may need to re-think their approaches to innovation altogether.

With this result, the leadership teams of AI-using organizations are again confronted with an important choice; even if AI could replace much of the human tasks in innovation, should it? From an efficiency perspective, automating as many tasks as possible may seem the obvious choice, freeing up resources for other work. Two arguments speak against it, however. Firstly, as already pointed out in Chapter 1, human actors may still be necessary to solve complex puzzles or complement algorithmic decision-making where codified pieces of information are missing. Removing these humans from the standard process of prediction may prevent them from learning the required skills, affecting the quality of both their prediction and judgement in more complex situations (Agrawal, Gans, & Goldfarb, 2019). Secondly, handing over control to algorithms could lead to unintended consequences affecting the well-being of society as a whole. Particularly the mechanism of framing, that is, creating demand for innovation by shaping consumer preferences, may be easily exploitable by algorithms possessing theory of mind – and thereby promise safer returns for lower investments than the comparatively uncertain research projects for actual breakthroughs. Allowing the algorithm to act freely and choose its most obviously promising strategy could lead to a dystopian future where organizations shift resources towards advertising-based exploitation rather than towards progress. Human control and oversight therefore continue to play essential roles in the process of AI-based innovative search.

Combined, Chapters 2, 3, and 4 thus strengthen the narrative emerging from Chapter 1; that managers must be acutely aware of the consequences of their actions, both positive and negative, in trying to navigate the introduction of artificial intelligence. Where they may feel the urge to intervene early and centralize control, they may be better advised to do the opposite; letting use case experimentation occur naturally (Chapter 2), and indeed continue to support it through strong investments in digitization (Chapter 3). To achieve deeper integration, managers may then again need to act against their intuitions, letting departments regain control after a period of centralization (Chapter 2), but continuing to invest in supporting structures and data availability (Chapter 3) – thereby creating a platform of supporting infrastructure. While short-term incentives to reduce costs by automating tasks are always strong, managers may instead need to possess the foresight to keep humans involved in routine and key decisions, thereby ensuring the build-up of the required capacities for expert judgement in difficult business- and moral choices (Chapter 4).

In short: Strategy matters. The choices made by managers and users of AI tools can significantly impact their organization’s success as we move forward in this AI age. In turn, as we define what AI means for organizations, we also define what it means for society. I believe that organizations can and must take on the responsibility of shaping policymaking – ensuring the ethical use of AI. Historic precedents show how key achievements of our modern world were introduced by responsibly acting business pioneers, before becoming institutionalized practice and law; for example the 5-day work week, introduced by Henry Ford in 1922, before being effectively singed into US law in 1938 (Blakemore, 2023) and making its way to Europe in only the 1960s (Futh, 2018). In a similar vein to such historic triumphs of societal well-being over business profits, leading businesses of today could set standards that shape the future of AI-related business practices; codes of conduct, for example regarding moral oversight of algorithmic work or the avoidance of labor replacement through task automation, can become lighthouses for peers and eventually regulated practice.

I thus close by reiterating the calls to action to three separate groups. Firstly, to academics: As outlined in the individual chapters of this dissertation, every question I attempted to answer leads to further important questions, promising to deepen our collective understanding of the drivers, opportunities, challenges, and effects of organizational AI use – and of organizational behavior in general. I call on academics to continue these conversations, maintaining a critical eye for the intended and unintended consequences AI may have. Secondly, to policymakers: To create an environment that not only allows organizations to reap the benefits of AI use and lets innovation flourish, but one that also safeguards the collective well-being, by preventing market failure, protecting personal privacy, and addressing the difficult question of how to fairly distribute the wealth created through technological innovation. And lastly, to managers: To pursue the paths that are best for their organizations, but to also reflect the broader ramifications of their choices, consider stakeholders as much as shareholders, and take seriously the responsibility this historic moment in time has placed in their hands.

Postscript:

To end this dissertation the way it began, I of course also tried letting AI write the final summary. I don't believe it quite works in getting the point across – just one small example of how human judgement still has a place in the AI age – but must concede that it is much more poetic than my version above:

```
-----------------------------------------------------------------
In this pivotal juncture of our technological evolution, the
utilization of artificial intelligence (AI) in organizational
settings presents both unprecedented opportunities and potential
pitfalls. As we conclude this dissertation, it is imperative to
issue a resounding call to action to the esteemed communities of
academics, policymakers, and managers. Academics, with their
scholarly rigor and intellectual acumen, are summoned to persist
in their pursuit of comprehensive insights into the multifaceted
applications of AI within organizations. Furthermore,
policymakers are implored to craft regulations that not only
facilitate the widespread accessibility of AI's benefits but also
safeguard against deleterious societal repercussions. Lastly,
managers, as the linchpins of operational implementation, are
urged to exercise ethical discernment in the deployment of AI,
thereby ensuring that its integration engenders positive
organizational outcomes while mitigating adverse social impacts.
It is incumbent upon us, as stewards of knowledge and
practitioners of management, to collectively steer the course of
AI's integration in organizations towards a future that is both
constructive and conscientious. Let us all, as scholars,
legislators, and business leaders, embrace the gravity of this
moment and act with the wisdom and integrity that the stewardship
of AI demands.

                                                   (GPT-4, 2024)[22].
-----------------------------------------------------------------
```

[22] Conversation with OpenAI's GPT-4, on May 20, 2024. Full text prompt used: "I am writing the final paragraph of my dissertation on the use of AI in organizations. I want to close with a call to action to three groups of actors: Academics, policymakers, and managers in this historic moment – to continue providing insights into AI use, shape regulation that allows for the benefits of AI use to reach the broad public, but avoid negative social consequences and use ethical judgement. Please draft this paragraph for me, using language appropriate for senior scholars in the field of management research."

LIST OF REFERENCES

Abada, I., & Lambin, X. (2023). Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided? *Management Science*, *69*(9), 5042-5065. https://doi.org/10.1287/mnsc.2022.4623

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. *NBER Working Paper No. w28257*. Retrieved 28 Apr 2023, from https://ssrn.com/abstract=3765910

Adner, R., & Levinthal, D. (2008). Doing versus seeing: acts of exploitation and perceptions of exploration. *Strategic Entrepreneurship Journal*, *2*(1), 43-52. https://doi.org/10.1002/sej.19

Adner, R., & Levinthal, D. A. (2004). What is not a real option: Considering boundaries for the application of real options to business strategy. *Academy of Management Review*, *29*(1), 74-85.

Agrawal, A., Gans, J., & Goldfarb, A. (2017). What to expect from artificial intelligence. *MIT Sloan Management Review*(13 July 2021). Retrieved 28 Apr 2023, from https://sloanreview.mit.edu/article/what-to-expect-from-artificial-intelligence/

Agrawal, A., Gans, J., & Goldfarb, A. (2019). Prediction, Judgement and Complexity. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 89-114). The University of Chicago Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2020). How to win with machine learning. *Harvard Business Review*, *2020*(Sep-Oct), 126-133.

Agrawal, A., Gans, J., & Goldfarb, A. (2022a). From Prediction to Transformation. *Harvard Business Review*, *2022*(Nov-Dec).

Agrawal, A., Gans, J., & Goldfarb, A. (2022b). *Power and Prediction - The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2022c). *Prediction Machines - The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.

Agrawal, A., Gans, J. S., & Goldfarb, A. (2024). Artificial intelligence adoption and system-wide change. *Journal of Economics & Management Strategy*, *33*(2), 327-337. https://doi.org/10.1111/jems.12521

Aguinis, H., Audretsch, D. B., Flammer, C., Meyer, K. E., Peng, M. W., & Teece, D. J. (2022). Bringing the Manager Back Into Management Scholarship. *Journal of Management*, *48*(7), 1849-1857. https://doi.org/10.1177/01492063221082555

Ahuja, G., Lampert, C. M., & Tandon, V. (2008). Moving beyond Schumpeter: Management research on the determinants of technological innovation. *Academy of Management Annals*, *2*(1), 1-98. https://doi.org/10.5465/19416520802211446

Allen, R., & Choudhury, P. (2022). Algorithm-Augmented Work and Domain Experience: The Countervailing Forces of Ability and Aversion. *Organization Science*, *33*(1), 149-169. https://doi.org/10.1287/orsc.2021.1554

Allen, T. (1970). Roles in Technical Communication Networks. In C. Nelson & D. Pollack (Eds.), *Communication Among Scientists and Engineers* (pp. 191-208). Heath Lexington Books.

Alvarez, S. A., & Porac, J. (2020). Imagination, indeterminacy, and managerial choice at the limit of knowledge. *Academy of Management Review*, *45*(4), 735-744. https://doi.org/10.5465/amr.2020.0366

Amabile, T. M. (1996). *Creativity and innovation in organizations* (Vol. 5). Harvard Business School Press.

Amabile, T. M. (2020). Creativity, Artificial Intelligence, and a World of Surprises [Guidepost]. *Academy of Management Discoveries*, *6*(3), 351-354. https://doi.org/10.5465/amd.2019.0075

Andriani, P., Ali, A., & Mastrogiorgio, M. (2017). Measuring Exaptation and Its Impact on Innovation, Search, and Problem Solving. *Organization Science*, *28*(2), 320-338. https://doi.org/10.1287/orsc.2017.1116

Andriani, P., & Cattani, G. (2016). Exaptation as source of creativity, innovation, and diversity: introduction to the Special Section. *Industrial and Corporate Change*, *25*(1), 115-131. https://doi.org/10.1093/icc/dtv053

Ansari, S., Garud, R., & Kumaraswamy, A. (2016). The disruptor's dilemma: TiVo and the U.S. television ecosystem. *Strategic Management Journal*, *37*(9), 1829-1853. https://doi.org/10.1002/smj.2442

Aral, S., Brynjolfsson, E., & Wu, L. (2012). Three-Way Complementarities: Performance Pay, Human Resource Analytics, and Information Technology. *Management Science*, *58*(5), 913-931. https://doi.org/10.1287/mnsc.1110.1460

Argote, L., & Greve, H. R. (2007). A Behavioral Theory of the Firm —40 Years and Counting: Introduction and Impact. *Organization Science*, *18*(3), 337–349. https://doi.org/10.1287/orsc.1070.0280

Athey, S., Bryan, K., & Gans, J. S. (2020). The Allocation of Decision Authority to Human and Artificial Intelligence. *NBER Working Paper No. w26673*. Retrieved 28 Apr 2023, from https://ssrn.com/abstract=3522322

Autor, D., & Salomons, A. (2018). Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share. *NBER Working Paper No. w24871*. Retrieved 8 Nov 2022, from https://ssrn.com/abstract=3222438

Bailey, D. E., Faraj, S., Hinds, P. J., Leonardi, P. M., & von Krogh, G. (2022). We Are All Theorists of Technology Now: A Relational Perspective on Emerging Technology and Organizing. *Organization Science*, *33*(1), 1-18. https://doi.org/10.1287/orsc.2021.1562
Balasubramanian, N., Ye, Y., & Xu, M. (2020). Substituting Human Decision-Making with Machine Learning: Implications for Organizational Learning. *Academy of Management Review*, *47*(3), 448-465. https://doi.org/10.5465/amr.2019.0470
Barley, S. R. (1986). Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly*, *31*(1), 78-108. https://doi.org/10.2307/2392767
Barley, S. R., & Kunda, G. (1992). Design and Devotion: Surges of Rational and Normative Ideologies of Control in Managerial Discourse. *Administrative Science Quarterly*, *37*(3), 363-399. https://doi.org/10.2307/2393449
Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning: Limitations and Opportunities. *fairmlbook.org*. Retrieved 13 May 2024, from https://fairmlbook.org/
Baron-Cohen, S. (1997). *Mindblindness: An essay on autism and theory of mind*. The MIT Press. https://doi.org/10.7551/mitpress/4635.001.0001
Beane, M. (2018). Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail. *Administrative Science Quarterly*, *64*(1), 87-123. https://doi.org/10.1177/0001839217751692
Beane, M., & Orlikowski, W. J. (2015). What Difference Does a Robot Make? The Material Enactment of Distributed Coordination. *Organization Science*, *26*(6), 1553-1573. https://doi.org/10.1287/orsc.2015.1004
Bell, J. J., Pescher, C., Tellis, G. J., & Füller, J. (2023). Can AI Help in Ideation? A Theory-Based Model for Idea Screening in Crowdsourcing Contests. *Marketing Science*, *43*(1), 54-72. https://doi.org/10.1287/mksc.2023.1434
Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., Darrell, T., Harari, Y. N., Zhang, Y.-Q., Xue, L., Shalev-Shwartz, S., Hadfield, G., Clune, J., Maharaj, T., Hutter, F., Baydin, A. G., McIlraith, S., Gao, Q., Acharya, A., Krueger, D., Dragan, A., Torr, P., Russell, S., Kahneman, D., Brauner, J., & Mindermann, S. (2024). Managing extreme AI risks amid rapid progress. *Science*, *384*(6698), 842-845. https://doi.org/10.1126/science.adn0117
Benner, M. J., & Tushman, M. L. (2003). Exploitation, Exploration, and Process Management: The Productivity Dilemma Revisited. *Academy of Management Review*, *28*(2), 238-256. https://doi.org/10.5465/amr.2003.9416096
Berg, J., Raj, M., & Seamans, R. (2023). Capturing Value from Artificial Intelligence. *Academy of Management Discoveries*. https://doi.org/10.5465/amd.2023.0106
Biernacki, P., & Waldorf, D. (1981). Snowball Sampling: Problems and Techniques of Chain Referral Sampling. *Sociological Methods & Research*, *10*(2), 141-163. https://doi.org/10.1177/004912418101000205
Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, *181*, 21-34. https://doi.org/10.1016/j.cognition.2018.08.003
Black, J. A., & Boal, K. B. (1996). Assessing the organizational capacity to change. In A. S. Heene, R. (Ed.), *Competence-based strategic management* (pp. 151-169). John Wiley & Sons.
Blakemore, E. (2023). How America settled on a 5-day workweek. *National Geographic*. Retrieved 15 May 2024, from https://www.nationalgeographic.com/history/article/american-workweek-history-explained
Bloom, N., Lemos, R., Sadun, R., Scur, D., & Van Reenen, J. (2016). International Data on Measuring Management Practices. *American Economic Review*, *106*(5), 152-156. https://doi.org/10.1257/aer.p20161058
Boal, K. B., & Hooijberg, R. (2001). Strategic leadership research: Moving on. *The Leadership Quarterly*, *11*(4), 515-549. https://doi.org/10.1016/S1048-9843(00)00057-6
Bock, L. (2015). *Work Rules!: Insights from Inside Google That Will Transform How You Live and Lead*. Twelve.
Bogost, I. (2022). ChatGPT Is Dumber Than You Think. *The Atlantic*. Retrieved 28 Apr 2023, from https://www.theatlantic.com/technology/archive/2022/12/chatgpt-openai-artificial-intelligence-writing-ethics/672386/
Bohm, C., Kirkpatrick, D., & Hintze, A. (2022). Understanding Memories of the Past in the Context of Different Complex Neural Network Architectures. *Neural Computation*, *34*(3), 754-780. https://doi.org/10.1162/neco_a_01469
Bostrom, N. (2016). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
Boumgarden, P., Nickerson, J., & Zenger, T. R. (2012). Sailing into the wind: Exploring the relationships among ambidexterity, vacillation, and organizational performance. *Strategic Management Journal*, *33*(6), 587-610. https://doi.org/10.1002/smj.1972
Bouschery, S. G., Blazevic, V., & Piller, F. T. (2023). Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management*, *40*(2), 139-153. https://doi.org/10.1111/jpim.12656

Bremner, R. P., & Eisenhardt, K. M. (2021). Organizing Form, Experimentation, and Performance: Innovation in the Nascent Civilian Drone Industry. *Organization Science*, *33*(4), 1645-1674. https://doi.org/10.1287/orsc.2021.1505

Bresnahan, T. (2021). Artificial intelligence technologies and aggregate growth prospects. In J. W. Diamond & G. R. Zodrow (Eds.), *Prospects for economic growth in the United States* (pp. 132-172). Cambridge University Press.

Bresnahan, T. F., & Trajtenberg, M. (1995). General Purpose Technologies 'Engines of Growth'? *Journal of Econometrics, 65*(1), 83-108. https://doi.org/10.1016/0304-4076(94)01598-T

Brock, J. K.-U., & Wangenheim, F. (2019). Demystifying AI: What digital transformation leaders can teach you about realistic artificial intelligence. *California Management Review*, *61*(4), 110-134.

Brynjolfsson, E. (1993). The Productivity Paradox of Information Technology. *Communications of the ACM*, *36*(12), 66-77.

Brynjolfsson, E., Jin, W., & McElheran, K. (2021). The power of prediction: Predictive analytics, workplace complements, and business performance. *Business Economics*, *56*(4), 217-239. https://doi.org/10.1057/s11369-021-00224-5

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530-1534. https://doi.org/10.1126/science.aap8062

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 23-60). The University of Chicago Press.

Brynjolfsson, E. T. W., & Hitt, L. (1996). Paradox lost? Firm-level evidence on the returns to information systems spending. *Management Science*, *42*(4), 541-555. https://doi.org/10.1287/mnsc.42.4.541

Brynjolfsson;, E., & McAfee, A. (2017). The Business of Artificial Intelligence. *Harvard Business Review*. Retrieved 28 Apr 2023, from https://hbr.org/2017/07/the-business-of-artificial-intelligence

Burns, T., & Stalker, G. M. (1961). *The Management of Innovation*. Tavistock.

Busch, C. (2024). Towards a Theory of Serendipity: A Systematic Review and Conceptualization. *Journal of Management Studies*, *61*(3), 1110-1151. https://doi.org/10.1111/joms.12890

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial. *Management Science*, *66*(2), 564-586. https://doi.org/10.1287/mnsc.2018.3249

Camuffo, A., Gambardella, A., & Pignataro, A. (2023). Framing Strategic Decisions in the Digital World. *Strategic Management Review*, *4*(2), 127-160. https://doi.org/10.1561/111.00000055

Candelon, F., Krayer, L., Rajendran, S., & Martínez, D. Z. (2023). How People Can Create - and Destroy - Value with Generative AI. *BCG Publications*. Retrieved 11 Nov 2023, from https://www.bcg.com/publications/2023/how-people-create-and-destroy-value-with-gen-ai

Carpenter, M. A. (2002). The implications of strategy and social context for the relationship between top management team heterogeneity and firm performance. *Strategic Management Journal*, *23*(3), 275-284. https://doi.org/10.1002/smj.226

Cattani, G., Sands, D., Porac, J., & Greenberg, J. (2018). Competitive Sensemaking in Value Creation and Capture. *Strategy Science*, *3*(4), 632-657. https://doi.org/10.1287/stsc.2018.0069

Certo, S. T., Busenbark, J. R., Woo, H. S., & Semadeni, M. (2016). Sample selection bias and Heckman models in strategic management research. *Strategic Management Journal*, *37*(13), 2639-2657. https://doi.org/10.1002/smj.2475

Chai, S. (2017). Near Misses in the Breakthrough Discovery Process. *Organization Science*, *28*(3), 411-428. https://doi.org/10.1287/orsc.2017.1134

Chakraa, H., Guérin, F., Leclercq, E., & Lefebvre, D. (2023). Optimization techniques for Multi-Robot Task Allocation problems: Review on the state-of-the-art. *Robotics and Autonomous Systems*, *168*. https://doi.org/10.1016/j.robot.2023.104492

Chan, H. C. S., Shan, H., Dahoun, T., Vogel, H., & Yuan, S. (2019). Advancing Drug Discovery via Artificial Intelligence. *Trends in Pharmacological Sciences*, *40*(8), 592-604. https://doi.org/10.1016/j.tips.2019.06.004

Chari, V. V., & Hopenhayn, H. (1991). Vintage Human Capital, Growth, and the Diffusion of New Technology. *Journal of Political Economy*, *99*(6), 1142-1165. https://doi.org/10.1086/261795

Choi, W. J., Liu, Q., & Shin, J. (2023). Predictive Analytics and Ship-Then-Shop Subscription. *Management Science*. https://doi.org/10.1287/mnsc.2023.4723

Chomsky, N., Roberts, I., & Watumull, J. (2023). The False Promise of ChatGPT. *New York Times*. Retrieved 13 Jun 2023, from https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

Chouard, T. (2016). The Go Files: AI computer wraps up 4-1 victory against human champion. *Nature*(15 March 2016). Retrieved 28 Apr 2023, from https://www.nature.com/articles/nature.2016.19575

Choudhary, V., Marchetti, A., Shrestha, Y. R., & Puranam, P. (2023). Human-AI Ensembles: When Can They Work? *Journal of Management*, 01492063231194968. https://doi.org/10.1177/01492063231194968
Choudhury, A. (2022). Toward an Ecologically Valid Conceptual Framework for the Use of Artificial Intelligence in Clinical Settings [Viewpoint]. *JMIR Human Factors*, *9*(2), e35421. https://doi.org/10.2196/35421
Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, *41*(8), 1381-1411. https://doi.org/10.1002/smj.3152
Clement, J., & Puranam, P. (2017). Searching for Structure: Formal Organization Design as a Guide to Network Evolution. *Management Science*, *64*(8), 3879-3895. https://doi.org/10.1287/mnsc.2017.2807
Colfer, L. J., & Baldwin, C. Y. (2016). The Mirroring Hypothesis: Theory, Evidence and Exceptions. *Industrial and Corporate Change*, *25*(5), 709-738. https://doi.org/10.1093/icc/dtw027
Colledanchise, M., & Ogren, P. (2016). How Behavior Trees Modularize Hybrid Control Systems and Generalize Sequential Behavior Compositions, the Subsumption Architecture and Decision Trees. *IEEE Transactions on Robotics*, *PP*. https://doi.org/10.1109/TRO.2016.2633567
Copeland, B. J. (2022). Artificial intelligence. *Encyclopedia Britannica*. Retrieved 28 Apr 2023, from https://www.britannica.com/technology/artificial-intelligence
Crafts, N. (2021). Artificial intelligence as a general-purpose technology: an historical perspective. *Oxford Review of Economic Policy*, *37*(3), 521-536. https://doi.org/10.1093/oxrep/grab012
Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.
Cronin, M. A., & Weingart, L. R. (2007). Representational gaps, information processing, and conflict in functionally diverse teams. *Academy of Management Review*, *32*(3), 761-773. https://doi.org/10.5465/amr.2007.25275511
Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. *Strategic Management Journal*, *37*(10), 2031-2049. https://doi.org/10.1002/smj.2440
Csaszar, F. A., & Steinberger, T. (2021). Organizations as Artificial Intelligences: The Use of Artificial Intelligence Analogies in Organization Theory. *Academy of Management Annals*, *16*(1), 1-37. https://doi.org/10.5465/annals.2020.0192
Cuzzolin, F., Morelli, A., Cîrstea, B., & Sahakian, B. J. (2020). Knowing me, knowing you: Theory of mind in AI. *Psychological Medicine*, *50*(7), 1057-1061. https://doi.org/10.1017/s0033291720000835
Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm* (2. ed.). Blackwell.
Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, *211*, 188-205. https://doi.org/10.1016/j.jebo.2023.05.008
Dahl, D. W., & Moreau, P. (2002). The Influence and Value of Analogical Thinking During New Product Ideation. *Journal of Marketing Research*, *39*(1), 47-60. https://doi.org/10.1509/jmkr.39.1.47.18930
Dash, R., McMurtrey, M., Rebman, C., & Kar, U. K. (2019). Application of Artificial Intelligence in Automation of Supply Chain Management. *Journal of Strategic Innovation & Sustainability*, *14*(3), 43-53. https://doi.org/10.33423/jsis.v14i3.2105
Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Retrieved 25 May 2024, from https://www.reuters.com/article/idUSKCN1MK0AG/
Davenport, T. H., & Mittal, N. (2022). How Generative AI Is Changing Creative Work. *Harvard Business Review*. Retrieved 28 Apr 2023, from https://hbr.org/2022/11/how-generative-ai-is-changing-creative-work
Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, *96*(1), 108-116.
Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. (2009). Optimal Structure, Market Dynamism, and the Strategy of Simple Rules. *Administrative Science Quarterly*, *54*(3), 413-452. https://doi.org/10.2189/asqu.2009.54.3.413
Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality *Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013*. Retrieved 09 Apr 2024, from https://ssrn.com/abstract=4573321
Denrell, J., Fang, C., & Liu, C. (2014). Perspective - Chance Explanations in the Management Sciences. *Organization Science*, *26*(3), 923-940. https://doi.org/10.1287/orsc.2014.0946
Denrell, J., Fang, C., & Liu, C. (2019). In Search of Behavioral Opportunities From Misattributions of Luck. *Academy of Management Review*, *44*(4), 896-915. https://doi.org/10.5465/amr.2017.0239
Department of Commerce, U. S. (2017). High-Tech Industries: The Role of FDI in Driving Innovation and Growth *SelectUSA*. Retrieved 15 Jan 2024, from https://www.trade.gov/sites/default/files/2021-10/High%20Tech%20Brief%20FINAL.pdf

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114-126. https://doi.org/10.1037/xge0000033

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science*, *64*(3), 1155-1170. https://doi.org/10.1287/mnsc.2016.2643

Dodgson, M., Ash, S., Andrews, J., & Phillips, N. (2022). Managing Technology-Enabled Innovation in a Professional Services Firm: A Cooperative Case Study. *Academy of Management Discoveries*, *8*(4), 509-530. https://doi.org/10.5465/amd.2020.0217

Donaldson, L. (2001). *The contingency theory of organizations*. Sage.

Doshi, A. R., Bell, J. J., Mirzayev, E., & Vanneste, B. (2024). Generative Artificial Intelligence and Evaluating Strategic Decisions. *SSRN*. Retrieved 15 May 2024, from https://ssrn.com/abstract=4714776

Edmondson, A. C., & McManus, S. E. (2007). Methodological fit in management field research. *Academy of Management Review*, *32*(4), 1246-1264. https://doi.org/10.5465/amr.2007.26586086

Eggers, J. P. (2014). Competing technologies and industry evolution: The benefits of making mistakes in the flat panel display industry. *Strategic Management Journal*, *35*(2), 159-178. https://doi.org/10.1002/smj.2129

Ehrig, T., & Foss, N. J. (2022). Why we need normative theories of entrepreneurial learning that go beyond Bayesianism. *Journal of Business Venturing Insights*, *18*, e00335. https://doi.org/10.1016/j.jbvi.2022.e00335

Eisenhardt, K. M. (1989). Building Theories from Case Study Research. *The Academy of Management Review*, *14*(4), 532. https://doi.org/10.2307/258557

Eisenhardt, K. M., & Bingham, C. B. (2017). Superior Strategy in Entrepreneurial Settings: Thinking, Doing, and the Logic of Opportunity. *Strategy Science*, *2*(4), 246-257. https://doi.org/10.1287/stsc.2017.0045

Eisenhardt, K. M., & Graebner, M. E. (2007). Theory Building From Cases: Opportunities And Challenges. *Academy of Management Journal*, *50*(1), 25-32. https://doi.org/10.5465/amj.2007.24160888

Eisenstein, M. (2021). Artificial intelligence powers protein-folding predictions. *Nature*, *599*, 706-708. Retrieved 28 Apr 2023, from https://www.nature.com/articles/d41586-021-03499-y

Ethiraj, S. K., & Levinthal, D. (2009). Hoping for A to Z While Rewarding Only A: Complex Organizations and Multiple Goals. *Organization Science*, *20*(1), 4-21. https://doi.org/10.1287/orsc.1080.0358

Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*, *21*(3), 625-642. https://doi.org/10.2307/40792435

Felin, T., Gambardella, A., Novelli, E., & Zenger, T. (2024). A Scientific Method for Startups. *Journal of Management*, 01492063231226136. https://doi.org/10.1177/01492063231226136

Felin, T., & Holweg, M. (2024). Theory Is All You Need: AI, Human Cognition, and Decision Making. *SSRN*. Retrieved 06 Mar 2024, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4737265

Felin, T., Koenderink, J., & Krueger, J. I. (2017). Rationality, perception, and the all-seeing eye. *Psychonomic Bulletin & Review*, *24*(4), 1040-1059. https://doi.org/10.3758/s13423-016-1198-z

Felin, T., & Zenger, T. R. (2009). Entrepreneurs as theorists. *Strategic Entrepreneurship Journal*, *3*(2), 127-146. https://doi.org/10.1002/sej.67

Felin, T., & Zenger, T. R. (2017). The Theory-Based View: Economic Actors as Theorists. *Strategy Science*, *2*(4), 258-271. https://doi.org/10.1287/stsc.2017.0048

Ferràs-Hernández, X. (2017). The Future of Management in a World of Electronic Brains. *Journal of Management Inquiry*, *27*(2), 260-263. https://doi.org/10.1177/1056492617724973

Finkelstein, S., & Hambrick, D. C. (1996). *Strategic leadership: Top executives and their effects on organizations*. West Publishing Company.

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). Principled Artificial Intelligence. *Berkman Klein Center for Internet & Society*(09 Dec 2021). Retrieved 28 Apr 2023, from http://nrs.harvard.edu/urn-3:HUL.InstRepos:42160420

Fleming, L. (2001). Recombinant Uncertainty in Technological Search. *Management Science*, *47*(1), 117-132. https://doi.org/10.1287/mnsc.47.1.117.10671

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

Foss, N. J., & Klein, P. G. (2012). *Organizing entrepreneurial judgment: A new approach to the firm*. Cambridge University Press.

Fountaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-Powered Organization. *Harvard Business Review*, *2019*(July-Aug), 62–73.

Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., Wang, D., Youn, H., & Rahwan, I. (2019). Toward understanding the impact of

artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, *116*(14), 6531-6539. https://doi.org/10.1073/pnas.1900949116
Freedman, D. H. (2010). Why Scientific Studies Are So Often Wrong: The Streetlight Effect. *Discover*. Retrieved 15 May 2024, from https://www.discovermagazine.com/the-sciences/why-scientific-studies-are-so-often-wrong-the-streetlight-effect
Futh, S. K. (2018). Die „Samstags gehört Vati mir"-Kampagne. In S. K. Futh (Ed.), *Strategische Kommunikation von Gewerkschaften* (pp. 147-195). Springer Fachmedien Wiesbaden. https://doi.org/10.1007/978-3-658-19583-0_5
Gaba, V., & Greve, H. R. (2019). Safe or Profitable? The Pursuit of Conflicting Goals. *Organization Science*, *30*(4), 647-667. https://doi.org/10.1287/orsc.2018.1280
Gadd, K. (2011). *TRIZ for Engineers: Enabling Inventive Problem Solving*. John Wiley & Sons.
Gaessler, F., & Piezunka, H. (2023). Training with AI: Evidence from chess computers. *Strategic Management Journal*, *44*(11), 2724-2750. https://doi.org/10.1002/smj.3512
Gambardella, A., Heaton, S., Novelli, E., & Teece, D. J. (2021). Profiting from Enabling Technologies? *Strategy Science*, *6*(1), 75-90. https://doi.org/10.1287/stsc.2020.0119
Gambardella, A., Khashabi, P., & Panico, C. (2020). Managing Autonomy in Industrial Research and Development: A Project-Level Investigation. *Organization Science*, *31*(1), 165-181. https://doi.org/10.1287/orsc.2019.1294
Gambardella, A., & McGahan, A. M. (2010). Business-Model Innovation: General Purpose Technologies and their Implications for Industry Structure. *Long Range Planning*, *43*(2–3), 262-271. https://doi.org/10.1016/j.lrp.2009.07.009
Gao, C., & McDonald, R. (2022). Shaping Nascent Industries: Innovation Strategy and Regulatory Uncertainty in Personal Genomics. *Administrative Science Quarterly*, *67*(4), 915-967. https://doi.org/10.1177/00018392221112641
Garisto, D. (2019). Google AI beats top human players at strategy game StarCraft II. *Nature*. Retrieved 28 Apr 2023, from https://www.nature.com/articles/d41586-019-03298-6/
Garud, R., Gehman, J., & Kumaraswamy, A. (2011). Complexity Arrangements for Sustained Innovation: Lessons from 3M Corporation. *Organization Studies*, *32*(6), 737-767. https://doi.org/10.1177/0170840611410810
Gates, B. (2023). The Age of AI has begun. *GatesNotes*. Retrieved 19 Jun 2023, from https://www.gatesnotes.com/The-Age-of-AI-Has-Begun
Gavetti, G. (2012). Toward a Behavioral Theory of Strategy. *Organization Science*, *23*(1), 267-285. https://doi.org/10.1287/orsc.1110.0644
Gavetti, G., Greve, H. R., Levinthal, D. A., & Ocasio, W. (2012). The Behavioral Theory of the Firm: Assessment and Prospects. *The Academy of Management Annals*, *6*(1), 1–40. https://doi.org/10.1080/19416520.2012.656841
Gavetti, G., Helfat, C. E., & Marengo, L. (2017). Searching, Shaping, and the Quest for Superior Performance. *Strategy Science*, *2*(3), 194-209. https://doi.org/10.1287/stsc.2017.0036
Gavetti, G., Levinthal, D. A., & Rivkin, J. W. (2005). Strategy Making in Novel and Complex Worlds: The Power of Analogy. *Strategic Management Journal*, *26*(8), 691-712. https://doi.org/10.2307/20142262
Gavetti, G., & Menon, A. (2016). Evolution Cum Agency: Toward a Model of Strategic Foresight. *Strategy Science*, *1*(3), 207-233. https://doi.org/10.1287/stsc.2016.0018
George, G. (2005). Slack Resources and the Performance of Privately Held Firms. *Academy of Management Journal*, *48*(4), 661-676. https://doi.org/10.5465/amj.2005.17843944
Georgieva, K. (2024). AI Will Transform the Global Economy. Let's Make Sure It Benefits Humanity. *International Monetary Fund*. Retrieved 16 May 2024, from https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity
Ghasemaghaei, M., Ebrahimi, S., & Hassanein, K. (2018). Data analytics competency for improving firm decision making performance. *The Journal of Strategic Information Systems*, *27*(1), 101-113. https://doi.org/10.1016/j.jsis.2017.10.001
Giermindl, L. M., Strich, F., Christ, O., Leicht-Deobald, U., & Redzepi, A. (2022). The dark sides of people analytics: reviewing the perils for organisations and employees. *European Journal of Information Systems*, *31*(3), 410-435. https://doi.org/10.1080/0960085X.2021.1927213
Gigerenzer, G. (2023). *The Intelligence of Intuition*. Cambridge University Press. https://doi.org/10.1017/9781009304887
Glaser, B., & Strauss, A. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Aldine.
Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals*, *14*(2), 627-660. https://doi.org/10.5465/annals.2018.0057
Goldfarb, A., Taska, B., & Teodoridis, F. (2023). Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings. *Research policy*, *52*(1). https://doi.org/10.1016/j.respol.2022.104653

Goodman, B., & Flaxman, S. (2017). European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation". *AI Magazine*, *38*(3), 50-57. https://doi.org/10.1609/aimag.v38i3.2741
GPT-4. (2024). *Text-based web chat on AI-related management research, accessed via You.com. Recipient: David M. Huber*. Retrieved 20 May 2024 from https://you.com/?chatMode=research
Graebner, M. E., Knott, A. M., Lieberman, M. B., & Mitchell, W. (2023). Empirical inquiry without hypotheses: A question-driven, phenomenon-based approach to strategic management research. *Strategic Management Journal*, *44*(1), 3-10. https://doi.org/10.1002/smj.3393
Grant, N. (2024). Google's A.I. Search Errors Cause a Furor Online. Retrieved 25 May 2024, from https://www.nytimes.com/2024/05/24/technology/google-ai-overview-search.html
Gregory, R. W., Henfridsson, O., Kaganer, E., & Kyriakou, H. (2021). The Role of Artificial Intelligence and Data Network Effects for Creating User Value. *Academy of Management Review*, *46*, 534-551. https://doi.org/10.5465/amr.2019.0178
Greve, H. R. (2003). A Behavioral Theory of R&D Expenditures and Innovations: Evidence from Shipbuilding. *Academy of Management Journal*, *46*(6), 685-702. https://doi.org/10.5465/30040661
Grodal, S., & O'Mahony, S. (2017). How does a Grand Challenge Become Displaced? Explaining the Duality of Field Mobilization. *Academy of Management Journal*, *60*(5), 1801-1827. https://doi.org/10.5465/amj.2015.0890
Grover, V., Chiang, R. H. L., Liang, T.-P., & Zhang, D. (2018). Creating Strategic Business Value from Big Data Analytics: A Research Framework. *Journal of Management Information Systems*, *35*(2), 388-423. https://doi.org/10.1080/07421222.2018.1451951
Haefner, N., Wincent, J., Parida, V., & Gassmann, O. (2021). Artificial intelligence and innovation management: A review, framework, and research agenda. *Technological Forecasting and Social Change*, *162*, 120392. https://doi.org/10.1016/j.techfore.2020.120392
Haenlein, M., & Kaplan, A. (2019). A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, *61*(4), 5-14. https://doi.org/10.1177/0008125619864925
Hagras, H. (2018). Toward Human-Understandable, Explainable AI. *Computer*, *51*(9), 28-36. https://doi.org/10.1109/MC.2018.3620965
Hambrick, D. C. (1989). Guest editor's introduction: Putting top managers back in the strategy picture. *Strategic Management Journal*, *10*(S1), 5-15. https://doi.org/10.1002/smj.4250100703
Hannah, D. P., & Eisenhardt, K. M. (2018). How firms navigate cooperation and competition in nascent ecosystems. *Strategic Management Journal*, *39*(12), 3163-3192. https://doi.org/10.1002/smj.2750
Hargadon, A. B., & Douglas, Y. (2001). When Innovations Meet Institutions: Edison and the Design of the Electric Light. *Administrative Science Quarterly*, *46*(3), 476-501. https://doi.org/10.2307/3094872
Hartmann, J., Exner, Y., & Domdey, S. (2023). The power of generative marketing: Can generative AI reach human-level visual marketing content? *SSRN*. Retrieved 15 May 2024, from https://ssrn.com/abstract=4597899
Hartmann, P., & Henkel, J. (2020). The Rise of Corporate Science in AI: Data as a Strategic Resource. *Academy of Management Discoveries*, *6*(3), 359-381. https://doi.org/10.5465/amd.2019.0043
Hasija, A., & Esper, T. L. (2022). In artificial intelligence (AI) we trust: A qualitative investigation of AI technology acceptance. *Journal of Business Logistics*, *43*(3), 388-412. https://doi.org/10.1111/jbl.12301
Hatchuel, A. (2002). Towards Design Theory and expandable rationality: The unfinished program of Herbert Simon. *Journal of Management and Governance*, *5*(3/4), 260-273. https://doi.org/10.1023/A:1014044305704
Hayward, M. L. A., Rindova, V. P., & Pollock, T. G. (2004). Believing one's own press: the causes and consequences of CEO celebrity. *Strategic Management Journal*, *25*(7), 637-653. https://doi.org/10.1002/smj.405
Heath, N. (2018). Microsoft to tackle AI skills shortage with two new training programs. *ZDNet*. Retrieved 28 Apr 2023, from https://www.zdnet.com/article/microsoft-to-tackle-ai-skills-shortage-with-two-new-training-programs/
Heaven, W. D. (2020). The way we train AI is fundamentally flawed. *MIT Technology Review*. Retrieved 28 Apr 2023, from https://www.technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-heath-nlp-computer-vision/
Heavey, C., Simsek, Z., Kyprianou, C., & Risius, M. (2020). How do strategic leaders engage with social media? A theoretical framework for research and practice. *Strategic Management Journal*, *41*(8), 1490-1527. https://doi.org/10.1002/smj.3156
Helfat, C. E. (2021). What Does Firm Shaping of Markets Really Mean? *Strategy Science*, *6*(4), 360-370. https://doi.org/10.1287/stsc.2021.0144
Heyes, C. (2015). Animal mindreading: what's the problem? *Psychonomic Bulletin & Review*, *22*(2), 313-327. https://doi.org/10.3758/s13423-014-0704-4

Hillman, A. (2021). *Bringing the Manager Back in Management*. Academy of Management. Retrieved 16 May 2024 from https://aom.org/events/annual-meeting/past-annual-meetings/2021-bringing-the-manager-back-in-management
Hindman, M. (2018). This is how Cambridge Analytica's Facebook targeting model really worked - according to the person who built it. Retrieved 25 May 2024, from https://www.niemanlab.org/2018/03/this-is-how-cambridge-analyticas-facebook-targeting-model-really-worked-according-to-the-person-who-built-it/
Hintze, A. (2016). Understanding the four types of AI, from reactive robots to self-aware beings. *The Conversation*. Retrieved 4 Apr 2024, from https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616
Hosmer, L. T. (1995). Trust: The Connecting Link Between Organizational Theory and Philosophical Ethics. *Academy of Management Review*, *20*(2), 379-403. https://doi.org/10.5465/amr.1995.9507312923
Hsieh, Y.-Y., & Vergne, J.-P. (2022). The future of the web? The coordination and early-stage growth of decentralized platforms. *Strategic Management Journal*. https://doi.org/10.1002/smj.3455
Huber, D. M. (2023). Introducing AI: The Search for Structure in Managing a General Purpose Technology. *Academy of Management Proceedings*, *2023*(1). https://doi.org/10.5465/AMPROC.2023.14804abstract
Huber, D. M. (2024, 15 July 2024). *Measuring AI Integration – A Quantification of Structural Determinants for AI Depth and Breadth in Organizations* Druid24 (accepted), Nice, France.
Huber, D. M., & Alexy, O. (2024). The Impact of Artificial Intelligence on Strategic Leadership. In Z. Simsek, C. Heavey, & B. Fox (Eds.), *Handbook of Research on Strategic Leadership in the Fourth Industrial Revolution* (Vol. in print). Edward Elgar Publishing.
Huber, D. M., & Reetz, D. K. (2024, 1 Mar 2024). *Organizing for AI - Multiple Goals, Structural Dynamics, and the Introduction of a General Purpose Technology* Organization Science Winter Conference, Zurich.
Huber, D. M., Reetz, D. K., & Alexy, O. (2021, 18 Nov 2021). *Innovation and the Asymptotic Rationality of Artificial Intelligence* Soph.I.A. Antipolis, Nice, France.
Huber, D. M., Reetz, D. K., & Alexy, O. (2022a, 20 Sep 2023). *Innovation and the Asymptotic Rationality of Artificial Intelligence* Strategic Management Society 42nd Annual Conference, London, UK.
Huber, D. M., Reetz, D. K., & Alexy, O. (2022b, 25 Jun 2022). *Innovation and the Asymptotic Rationality of Artificial Intelligence* Vienna Conference on Strategy, Organizational Design and Innovation, Vienna.
Hyatt, R. (1981). Checkmate - an electronic chess champion. *Cray Channels*, *3*(2), 6-11. Retrieved 09 Jan 2023, from http://archive.computerhistory.org/projects/chess/related_materials/text/3-1%20and%203-2.Cray_Channels_Vol-3_No-2.Checkmate.Cray_Blitz.Hyatt.1981/Cray_Channels_Vol-3_No-2.Checkmate.Cray_Blitz.Hyatt.1981.062303019.sm.pdf
Iansiti, M., & Lakhani, K. R. (2020). *Competing in the age of AI : strategy and leadership when algorithms and networks run the world*. Harvard Business Review Press.
Iansiti, M., & Nadella, S. (2022). Democratizing Transformation. *Harvard Business Review*, *May-June*(2022). Retrieved 19 Dec 2022, from https://hbr.org/2022/05/democratizing-transformation
IBM. (2022a). Explainable AI. *Watson*. Retrieved 28 Apr 2023, from https://www.ibm.com/watson/explainable-ai
IBM. (2022b). IBM Global AI Adoption Index 2022. *Watson*. Retrieved 27 Feb 2024, from https://www.ibm.com/downloads/cas/GVAGA3JP
Ichniowski, C., Shaw, K., & Prennushi, G. (1997). The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines. *The American Economic Review*, *87*(3), 291-313.
Jaffe, A. B. (1986). Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value. *American Economic Review*, *76*(5), 984-1001.
Janssen, M., van der Voort, H., & Wahyudi, A. (2017). Factors influencing big data decision-making quality. *Journal of Business Research*, *70*, 338-345. https://doi.org/10.1016/j.jbusres.2016.08.007
Jia, N., Luo, X., Fang, Z., & Liao, C. (2023). When and How Artificial Intelligence Augments Employee Creativity. *Academy of Management Journal*. https://doi.org/10.5465/amj.2022.0426
Joseph, J., & Gaba, V. (2020). Organizational Structure, Information Processing, and Decision-Making: A Retrospective and Road Map for Research. *Academy of Management Annals*, *14*(1), 267-302. https://doi.org/10.5465/annals.2017.0103
Jovanovic, B., & Rousseau, P. L. (2005). General Purpose Technologies. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of Economic Growth* (Vol. 1B, pp. 1181-1224). Elsevier. https://doi.org/10.1016/S1574-0684(05)01018-X
Just, J., Ströhle, T., Füller, J., & Hutter, K. (2023). AI-based novelty detection in crowdsourced idea spaces. *Innovation*, 1-28. https://doi.org/10.1080/14479338.2023.2215740
Katila, R., & Ahuja, G. (2002). Something Old, Something New: A Longitudinal Study of Search Behavior and New Product Introduction. *The Academy of Management Journal*, *45*(6), 1183-1194. https://doi.org/10.2307/3069433
Kawaguchi, K. (2020). When Will Workers Follow an Algorithm? A Field Experiment with a Retail Business. *Management Science*, *67*(3), 1670-1695. https://doi.org/10.1287/mnsc.2020.3599

Kearney, J. (2019). A Critical Analysis of the Applications and Limitations of Organizational Network Analysis in the Workplace to Drive Organizational Change. *SSRN*. Retrieved 28 Apr 2023, from https://ssrn.com/abstract=3564195

Kelion, L. (2020). DeepMind's AI agent MuZero could turbocharge YouTube. *BBC News*(23 Dec 2020). Retrieved 09 Jan 2023, from https://www.bbc.com/news/technology-55403473

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at Work: The New Contested Terrain of Control. *Academy of Management Annals*, *14*(1), 366-410. https://doi.org/10.5465/annals.2018.0174

Khatib, F., DiMaio, F., Cooper, S., Kazmierczyk, M., Gilski, M., Krzywda, S., Zabranska, H., Pichova, I., Thompson, J., Popović, Z., Jaskolski, M., Baker, D., Foldit Contenders, G., & Foldit Void Crushers, G. (2011). Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature Structural & Molecular Biology*, *18*(10), 1175-1177. https://doi.org/10.1038/nsmb.2119

Kim, S., Wang, Y., & Boon, C. (2021). Sixty years of research on technology and human resource management: Looking back and looking forward. *Human Resource Management*, *60*(1), 229-247. https://doi.org/10.1002/hrm.22049

Kim, S. Y., & Upneja, A. (2021). Majority voting ensemble with a decision trees for business failure prediction during economic downturns. *Journal of Innovation & Knowledge*, *6*(2), 112-123. https://doi.org/10.1016/j.jik.2021.01.001

Kiron, D., & Schrage, M. (2019). Strategy for and with AI. *MIT Sloan Management Review*, *60*(4), 29-36.

Kittur, A., Yu, L., Hope, T., Chan, J., Lifshitz-Assaf, H., Gilon, K., Ng, F., Kraut, R. E., & Shahaf, D. (2019). Scaling up analogical innovation with crowds and AI. *Proceedings of the National Academy of Sciences*, *116*(6), 1870-1877. https://doi.org/10.1073/pnas.1807185116

Kneeland, M. K., Schilling, M. A., & Aharonson, B. S. (2020). Exploring Uncharted Territory: Knowledge Search Processes in the Origination of Outlier Innovation. *Organization Science*, *31*(3), 535-557. https://doi.org/10.1287/orsc.2019.1328

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin.

Knight, W. (2022). ChatGPT's Most Charming Trick Is Also Its Biggest Flaw. *Wired*. Retrieved 9 Jan 2023, from https://www.wired.com/story/openai-chatgpts-most-charming-trick-hides-its-biggest-flaw/

Knudsen, T., & Srikanth, K. (2014). Coordinated Exploration: Organizing Joint Search by Multiple Specialists to Overcome Mutual Confusion and Joint Myopia. *Administrative Science Quarterly*, *59*(3), 409-441. https://doi.org/10.1177/0001839214538021

Koçak, Ö., Levinthal, D. A., & Puranam, P. (2022). The Dual Challenge of Search and Coordination for Organizational Adaptation: How Structures of Influence Matter. *Organization Science*, *34*(2), 851-869. https://doi.org/10.1287/orsc.2022.1601

Kosinski, M. (2023). Evaluating Large Language Models in Theory of Mind Tasks. *ArXiv*. Retrieved 21 Mar 2024, from https://arxiv.org/abs/2302.02083

Krakowski, S., Luger, J., & Raisch, S. (2022). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, *44*(6), 1425–1452. https://doi.org/10.1002/smj.3387

Kretschmer, T., & Khashabi, P. (2020). Digital Transformation and Organization Design: An Integrated Approach. *California Management Review*, *62*(4), 86-104. https://doi.org/10.1177/0008125620940296

Kretschmer, T., Leiponen, A., Schilling, M., & Vasudeva, G. (2022). Platform ecosystems as meta-organizations: Implications for platform strategies. *Strategic Management Journal*, *43*(3), 405-424. https://doi.org/10.1002/smj.3250

Kruhse-Lehtonen, U., & Hofmann, D. (2020). How to Define and Execute Your Data and AI Strategy. *Harvard Data Science Review*, *2*(3). Retrieved 28 Apr 2023, from https://hdsr.mitpress.mit.edu/pub/4vlrf0x2/release/2

Langley, A. (1999). Strategies for Theorizing from Process Data. *Academy of Management Review*, *24*(4), 691-710. https://doi.org/10.5465/amr.1999.2553248

Lantz, B. (2019). *Machine Learning with R* (3rd ed.). Packt Publishing.

Lanzolla, G., Lorenz, A., Miron-Spektor, E., Schilling, M., Solinas, G., & Tucci, C. L. (2020). Digital transformation: What is new if anything? Emerging patterns and management research. *6*, 341-350. https://doi.org/10.5465/amd.2020.0144

Le Masson, P., Hatchuel, A., Le Glatin, M., & Weil, B. (2019). Designing Decisions in the Unknown: A Generative Model. *European Management Review*, *16*(2), 471-490. https://doi.org/10.1111/emre.12289

Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021). Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What. *MIS Quarterly*, *45*(3), 1501-1526. https://doi.org/10.25300/MISQ/2021/16564

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science*, *33*(1), 126-148. https://doi.org/10.1287/orsc.2021.1549

Lee, K.-F. (2018). *AI Superpowers: China, Silicon Valley, and the New World Order*. Houghton Mifflin Harcourt.

Leonardi, P. (2020). You're going digital-now what. *MIT Sloan Management Review*, *61*(2), 28-35.

Levinthal, D. A. (1997). Adaptation on Rugged Landscapes. *Management Science*, *43*(7), 934-950.
Levinthal, D. A. (2011). A behavioral approach to strategy - what's the alternative? *Strategic Management Journal*, *32*(13), 1517-1523. https://doi.org/10.1002/smj.963
Levinthal, D. A., & Rerup, C. (2020). The Plural of Goal: Learning in a World of Ambiguity. *Organization Science*, *32*(3), 527-543. https://doi.org/10.1287/orsc.2020.1383
Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *ArXiv*. Retrieved 21 Mar 2024, from https://arxiv.org/abs/2005.11401v4
Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. (2020). Explainable AI: A review of machine learning interpretability methods. *Entropy*, *23*(1), 18. https://doi.org/10.3390/e23010018
Liu, D., Fisher, G., & Chen, G. (2018). CEO Attributes and Firm Performance: A Sequential Mediation Process Model. *Academy of Management Annals*, *12*(2), 789-816. https://doi.org/10.5465/annals.2016.0031
Liu, Y. Z., Qian, Y., Jiang, Y. C., & Shang, J. (2020). Using favorite data to analyze asymmetric competition: Machine learning models. *European Journal of Operational Research*, *287*(2), 600-615. https://doi.org/10.1016/j.ejor.2020.03.074
Llama3. (2024). *Text-based web chat on AI-related management research, accessed via You.com. Recipient: David M. Huber*. Retrieved 10 May 2024 from https://you.com/?chatMode=research
Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90-103. https://doi.org/10.1016/j.obhdp.2018.12.005
Lu, J., Lee, D., Kim, T. W., & Danks, D. (2019). Good Explanation for Algorithmic Transparency. *SSRN*. Retrieved 28 Apr 2023, from https://papers.ssrn.com/abstract_id=3503603
Ludwig, J., & Mullainathan, S. (2023). Machine Learning as a Tool for Hypothesis Generation. *NBER Working Paper*. Retrieved 15 Feb 2024, from https://www.nber.org/papers/w31017
MacAulay, S. C., Steen, J., & Kastelle, T. (2020). The search environment is not (always) benign: reassessing the risks of organizational search. *Industrial and Corporate Change*, *29*(1), 1-23. https://doi.org/10.1093/icc/dtx045
March, J. G. (1962). The business firm as a political coalition. *The Journal of politics*, *24*(4), 662-678.
March, J. G. (1978). Bounded Rationality, Ambiguity, and the Engineering of Choice. *The Bell Journal of Economics*, *9*(2), 587-608. https://doi.org/10.2307/3003600
March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, *2*(1), 71-87. https://doi.org/10.1287/orsc.2.1.71
March, J. G. (2006). Rationality, foolishness, and adaptive intelligence. *Strategic Management Journal*, *27*(3), 201-214. https://doi.org/10.1002/smj.515
March, J. G., & Simon, H. A. (1958). *Organizations*. Wiley.
March, J. G., & Simon, H. A. (1993). *Organizations* (2nd ed.). Wiley.
Markus, M. L., & Robey, D. (1988). Information Technology and Organizational Change: Causal Structure in Theory and Research. *Management Science*, *34*(5), 583-598. https://doi.org/10.1287/mnsc.34.5.583
Marr, B. (2023). The Top 10 Limitations Of ChatGPT. *Forbes*. Retrieved 28 Apr 2022, from https://www.forbes.com/sites/bernardmarr/2023/03/03/the-top-10-limitations-of-chatgpt/
Martin, K. E. (2019). Designing ethical algorithms. *MIS Quarterly Executive June*, *18*(2), 129–142.
McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. *Stanford Engineering*. Retrieved 01 Apr 2024, from http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf
McCorduck, P. (2004). *Machines Who Think* (2nd ed.). A.K.Peters.
McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI Adoption in America: Who, What, and Where. *Journal of Economics & Management Strategy*, *33*(2), 375-415. https://doi.org/10.1111/jems.12576
Mee, J. F. (2022). Frederick W. Taylor. *Encyclopedia Britannica*. Retrieved 28 Apr 2023, from https://www.britannica.com/biography/Frederick-W-Taylor
Mercedes-Benz. (2022). Two Letters and Four Principles: How Mercedes-Benz Uses Artificial Intelligence (AI). *Sustainability*. Retrieved 28 Apr 2023, from https://group.mercedes-benz.com/sustainability/data/ki-guidelines.html
Merendino, A., Dibb, S., Meadows, M., Quinn, L., Wilson, D., Simkin, L., & Canhoto, A. (2018). Big Data, Big Decisions: The Impact of Big Data on Board Level Decision-Making. *Journal of Business Research*, *93*, 67-78. https://doi.org/10.1016/j.jbusres.2018.08.029
Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An Expanded Sourcebook* (2nd ed.). Sage Publications.
Miller, D. (1996). Configurations Revisited. *Strategic Management Journal*, *17*(7), 505-512. https://doi.org/10.1002/(SICI)1097-0266(199607)17:7<505::AID-SMJ852>3.0.CO;2-I
Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The Structure of "Unstructured" Decision Processes. *Administrative Science Quarterly*, *21*(2), 246-275. https://doi.org/10.2307/2392045

Mitchell, M. (2019). Artificial Intelligence Hits the Barrier of Meaning. *Information*, *10*(2), 51. https://doi.org/10.3390/info10020051
Mohamad Suhaili, S., Salim, N., & Jambli, M. N. (2021). Service chatbots: A systematic review. *Expert Systems with Applications*, *184*, 115461. https://doi.org/10.1016/j.eswa.2021.115461
Mollick, E. (2024). *Co-Intelligence*. WH Allen.
Mount, M. P., Ertug, G., Kavusan, K., George, G., & Zou, T. (2024). Reeling in the Slack: An Integrative Review to Reinstate Slack as a Central Theoretical Construct for Management Research. *Academy of Management Annals*, *0*(0), 1-33. https://doi.org/10.5465/annals.2023.0087
Murray, A., Rhymer, J., & Sirmon, D. G. (2021). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review*, *46*(3), 552-571. https://doi.org/10.5465/amr.2019.0186
Nelson, R. R., & Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.
Newell, A., Shaw, J. C., & Simon, H. A. (1957). Empirical explorations of the Logic Theory Machine: A case study in heuristics. *Proceedings of the 1957 Western Joint Computer Conference*, 218-230.
Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Prentice-Hall.
Nickerson, J., Yen, C. J., & Mahoney, J. T. (2012). Exploring the Problem-Finding and Problem-Solving Approach for Designing Organizations. *Academy of Management Perspectives*, *26*(1), 52-72. https://doi.org/10.5465/amp.2011.0106
Nickerson, J. A., & Zenger, T. R. (2004). A Knowledge-Based Theory of the Firm: The Problem-Solving Perspective. *Organization Science*, *15*(6), 617-632.
Nohria, N., & Gulati, R. (1996). Is Slack Good or Bad for Innovation? *Academy of Management Journal*, *39*(5), 1245-1264. https://doi.org/10.5465/256998
Norvik, P., & Russell, S. (2021). *Artificial Intelligence: A Modern Approach* (4th Global ed.). Pearson Education Ltd.
OECD. (2015). The R&D Intensity of Economic Activities in OECD Countries. *OECD Economic Analysis and Statistics Division (EAS)*. Retrieved 27 Nov 2023, from https://one.oecd.org/document/DSTI/EAS/STP/NESTI(2015)9/en/pdf
Ott, T. E., & Eisenhardt, K. M. (2020). Decision weaving: Forming novel, complex strategy in entrepreneurial settings. *Strategic Management Journal*, *41*(12), 2275-2314. https://doi.org/10.1002/smj.3189
Ozcan, P., & Eisenhardt, K. M. (2009). Origin of Alliance Portfolios: Entrepreneurs, Network Strategies, and Firm Performance. *Academy of Management Journal*, *52*(2), 246-279. https://doi.org/10.5465/amj.2009.37308021
Pachidi, S., Berends, H., Faraj, S., & Huysman, M. (2021). Make Way for the Algorithms: Symbolic Actions and Change in a Regime of Knowing. *Organization Science*, *32*(1), 18-41. https://doi.org/10.1287/orsc.2020.1377
Packard, M. D., & Clark, B. B. (2020). On the Mitigability of Uncertainty and the Choice between Predictive and Nonpredictive Strategy. *Academy of Management Review*, *45*(4), 766-786. https://doi.org/10.5465/amr.2018.0198
Park, M., Leahey, E., & Funk, R. J. (2023). Papers and patents are becoming less disruptive over time. *Nature*, *613*(7942), 138-144. https://doi.org/10.1038/s41586-022-05543-x
Parmar, B., & Freeman, R. E. (2016). Ethics and the algorithm. *MIT Sloan Management Review*, *58*(1), 10-11.
Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books.
Peirce, C. S. (1957). The logic of abduction. In V. Thomas (Ed.), *Peirce's essays in the philosophy of science* (pp. 195-205). Liberal Arts Press.
Posen, H. E., Keil, T., Kim, S., & Meissner, F. D. (2018). Renewing Research on Problemistic Search - A Review and Research Agenda. *Academy of Management Annals*, *12*(1), 208-251. https://doi.org/10.5465/annals.2016.0018
Posen, H. E., Leiblein, M. J., & Chen, J. S. (2018). Toward a behavioral theory of real options: Noisy signals, bias, and learning. *Strategic Management Journal*, *39*(4), 1112-1138. https://doi.org/10.1002/smj.2757
Posen, H. E., & Levinthal, D. A. (2012). Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments. *Management Science*, *58*(3), 587-601. https://doi.org/10.1287/mnsc.1110.1420
Prahl, A., & Van Swol, L. (2017). Understanding algorithm aversion: When is advice from automation discounted? *Journal of Forecasting*, *36*(6), 691-702. https://doi.org/10.1002/for.2464
Pratt, M. G., Sonenshein, S., & Feldman, M. S. (2020). Moving Beyond Templates: A Bricolage Approach to Conducting Trustworthy Qualitative Research. *Organizational Research Methods*, *25*(2), 211-238. https://doi.org/10.1177/1094428120927466
Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, *1*(4), 515-526. https://doi.org/10.1017/S0140525X00076512
Puranam, P. (2020). Human-AI Collaborative Decision Making as an Organization Design Problem. *INSEAD Working Paper 2020/50/STR*. Retrieved 28 Apr 2023, from https://ssrn.com/abstract=3717500
Puranam, P., Alexy, O., & Reitzig, M. (2014). What's "New" About New Forms of Organizing? *Academy of Management Review*, *39*(2), 162–180. https://doi.org/10.5465/amr.2011.0436

Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). Modelling Bounded Rationality in Organizations: Progress and Prospects. *Academy of Management Annals*, *9*(1), 337-392. https://doi.org/10.5465/19416520.2015.1024498
Raisch, S., & Fomina, K. (2021). Shifting the Bounds on Rationality: Artificial Intelligence and Organizational Search. *Academy of Management Proceedings*, *2021*(1), 13400. https://doi.org/10.5465/AMBPP.2021.13400abstract
Raisch, S., & Fomina, K. (2024). Combining Human and Artificial Intelligence: Hybrid Problem-Solving in Organizations. *Academy of Management Review*, *0*(0). https://doi.org/10.5465/amr.2021.0421
Raisch, S., & Krakowski, S. (2020). Artificial Intelligence and Management: The Automation-Augmentation Paradox. *Academy of Management Review*. https://doi.org/10.5465/2018.0072
Raj, M., & Seamans, R. (2019). Primer on artificial intelligence and robotics. *Journal of Organization Design*, *8*(1), 11. https://doi.org/10.1186/s41469-019-0050-0
Rammer, C., Fernández, G. P., & Czarnitzki, D. (2022). Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Research policy*, *51*(7), 104555. https://doi.org/10.1016/j.respol.2022.104555
Rindova, V., & Courtney, H. (2020). To Shape or Adapt: Knowledge Problems, Epistemologies, and Strategic Postures under Knightian Uncertainty. *Academy of Management Review*, *45*(4), 787-807. https://doi.org/10.5465/amr.2018.0291
Rindova, V. P., & Kotha, S. (2001). Continuous "Morphing": Competing through Dynamic Capabilities, Form, and Function. *Academy of Management Journal*, *44*(6), 1263-1280. https://doi.org/10.5465/3069400
Rindova, V. P., & Martins, L. L. (2017). From Values to Value: Value Rationality and the Creation of Great Strategies. *Strategy Science*, *3*(1), 323-334. https://doi.org/10.1287/stsc.2017.0038
Rindova, V. P., & Martins, L. L. (2021). Shaping Possibilities: A Design Science Approach to Developing Novel Strategies. *Academy of Management Review*, *46*(4), 800-822. https://doi.org/10.5465/amr.2019.0289
Rindova, V. P., & Martins, L. L. (2023). Moral Imagination, the Collective Desirable, and Strategic Purpose. *Strategy Science*, *8*(2), 170-181. https://doi.org/10.1287/stsc.2023.0190
Rittel, H. W., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, *4*, 155-169.
Rivkin, J. W., & Siggelkow, N. (2006). Organizing to Strategize in the Face of Interactions: Preventing Premature Lock-in. *Long Range Planning*, *39*(6), 591-614. https://doi.org/10.1016/j.lrp.2006.10.005
Roach, J. (2018). Aiming to fill skill gaps in AI, Microsoft makes training courses available to the public. *Microsoft: The AI Blog*. Retrieved 28 Apr 2023, from https://blogs.microsoft.com/ai/microsoft-professional-program-ai/
Robey, D., & Boudreau, M.-C. (1999). Accounting for the Contradictory Organizational Consequences of Information Technology: Theoretical Directions and Methodological Implications. *Information Systems Research*, *10*(2), 167-185. https://doi.org/10.1287/isre.10.2.167
Rosenkopf, L., & Nerkar, A. (2001). Beyond Local Search: Boundary-spanning, Exploration, and Impact in the Optical Disk Industry. *Strategic Management Journal*, *22*(4), 287-306. https://doi.org/10.1002/smj.160
Samimi, M., Cortes, A. F., Anderson, M. H., & Herrmann, P. (2022). What is strategic leadership? Developing a framework for future research. *The Leadership Quarterly*, *33*(3), 101353. https://doi.org/10.1016/j.leaqua.2019.101353
Sap, M., Le Bras, R., Fried, D., & Choi, Y. (2022). Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs. *ArXiv*. Retrieved 15 Apr 2024, from https://arxiv.org/abs/2210.13312
Satariano, A., & Kang, C. (2023). How Nations Are Losing a Global Race to Tackle A.I.'s Harms. *The New York Times*. Retrieved 15 May 2024, from https://www.nytimes.com/2023/12/06/technology/ai-regulation-policies.html
Schafheitle, S., Weibel, A., Ebert, I., Kasper, G., Schank, C., & Leicht-Deobald, U. (2020). No Stone Left Unturned? Toward a Framework for the Impact of Datafication Technologies on Organizational Control. *Academy of Management Discoveries*, *6*(3), 455-487. https://doi.org/10.5465/amd.2019.0002
Schendel, D., & Hofer, C. W. ((eds.). 1979). *Strategic management: A new view of business policy and planning*. Little, Brown.
Schilling, M. (2018). *Quirky: The remarkable story of the traits, foibles, and genius of breakthrough innovators who changed the world*. PublicAffairs.
Schilling, M. A. (2017). The Cognitive Foundations of Visionary Strategy. *Strategy Science*, *3*(1), 335-342. https://doi.org/10.1287/stsc.2017.0044
Schossau, J., & Hintze, A. (2023). *Towards a Theory of Mind for Artificial Intelligence Agents* ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference, https://doi.org/10.1162/isal_a_00605
Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T., & Silver, D. (2020a). Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, *588*(7839), 604-609. https://doi.org/10.1038/s41586-020-03051-4

Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T., & Silver, D. (2020b). MuZero: Mastering Go, chess, shogi and Atari without rules. *DeepMind Blog*(23 Dec 2020). Retrieved 09 Jan 2023, from https://deepmind.com/blog/article/muzero-mastering-go-chess-shogi-and-atari-without-rules

Scur, D., Sadun, R., Van Reenen, J., Lemos, R., & Bloom, N. (2021). The World Management Survey at 18: lessons and the way forward. *Oxford Review of Economic Policy*, *37*(2), 231-258. https://doi.org/10.1093/oxrep/grab009

Seppälä, J. (2021). AI Affecting Local Search in Strategic Decision-Making. *Academy of Management Proceedings*, *2021*(1), 14630. https://doi.org/10.5465/AMBPP.2020.14630abstract

Sergeeva, A. V., Faraj, S., & Huysman, M. (2020). Losing Touch: An Embodiment Perspective on Coordination in Robotic Surgery. *Organization Science*, *31*(5), 1248-1271. https://doi.org/10.1287/orsc.2019.1343

Shanahan, M. (2016). The Frame Problem. *The Stanford Encyclopedia of Philosophy*. Retrieved 28 Apr 2023, from https://plato.stanford.edu/archives/spr2016/entries/frame-problem/

Shepardson, D., Warburton, M., & Stone, M. (2023). Tech titans meet US lawmakers, Musk seeks 'referee' for AI. *Reuters*. Retrieved 16 May 2024, from https://www.reuters.com/technology/musk-zuckerberg-gates-join-us-senators-ai-forum-2023-09-13/

Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational Decision-Making Structures in the Age of Artificial Intelligence. *California Management Review*, *61*(4), 66-83. https://doi.org/10.1177/0008125619862257

Shrestha, Y. R., He, V. F., Puranam, P., & von Krogh, G. (2020). Algorithm Supported Induction for Building Theory: How Can We Use Prediction Models to Theorize? *Organization Science*, *32*(3), 856-880. https://doi.org/10.1287/orsc.2020.1382

Shrestha, Y. R., He, V. F., Puranam, P., & von Krogh, G. (2021). Algorithm Supported Induction for Building Theory: How Can We Use Prediction Models to Theorize? *Organization Science*, *32*(3), 856-880. https://doi.org/10.1287/orsc.2020.1382

Siggelkow, N. (2002). Evolution toward Fit. *Administrative Science Quarterly*, *47*(1), 125-159. https://doi.org/10.2307/3094893

Siggelkow, N., & Levinthal, D. A. (2003). Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation. *Organization Science*, *14*(6), 650-669. https://doi.org/10.1287/orsc.14.6.650.24840

Siggelkow, N., & Rivkin, J. W. (2005). Speed and Search: Designing Organizations for Turbulence and Complexity. *Organization Science*, *16*(2), 101-122. https://doi.org/10.1287/orsc.1050.0116

Siggelkow, N., & Rivkin, J. W. (2009). Hiding the Evidence of Valid Theories: How Coupled Search Processes Obscure Performance Differences among Organizations. *Administrative Science Quarterly*, *54*(4), 602-634. https://doi.org/10.2189/asqu.2009.54.4.602

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, *69*(1), 99-118.

Simon, H. A. (1957). *Administrative Behavior* (2 ed.). Macmillan.

Simon, H. A. (1964). On the Concept of Organizational Goal. *Administrative Science Quarterly*, *9*(1), 1-22. https://doi.org/10.2307/2391519

Simon, H. A. (1987). Making Management Decisions: the Role of Intuition and Emotion. *Academy of Management Perspectives*, *1*(1), 57-64. https://doi.org/10.5465/ame.1987.4275905

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*(1), 1-20.

Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.). MIT Press. (1969)

Simon, H. A. (1997). *Administrative Behavior* (Fourth ed.). The Free Press. (1947)

Simsek, Z., Heavey, C., & Fox, B. C. (2021). A managerial interfaces perspective for competitive action research. *Long Range Planning*, *102141*. https://doi.org/10.1016/j.lrp.2021.102141

Singh, R., Sledzieski, S., Bryson, B., Cowen, L., & Berger, B. (2023). Contrastive learning in protein language space predicts interactions between drugs and protein targets. *Proceedings of the National Academy of Sciences*, *120*(24), e2220778120. https://doi.org/10.1073/pnas.2220778120

Smith, P., & Beretta, M. (2021). The gordian knot of practicing digital transformation: coping with emergent paradoxes in ambidextrous organizing structures. *Journal of Product Innovation Management*, *38*(1), 166-191. https://doi.org/10.1111/jpim.12548

Smith, W. K., & Lewis, M. W. (2011). Toward a Theory of Paradox: A Dynamic equilibrium Model of Organizing. *Academy of Management Review*, *36*(2), 381-403. https://doi.org/10.5465/amr.2009.0223

Spradley, J. P. (1979). *The Ethnographic Interview*. Holt, Rinehart and Winston.

Strauss, A. L., & Corbin, J. (1990). *Basics of qualitative research: Grounded theory procedures and techniques*. Sage.

Suddaby, R. (2006). From the Editors: What Grounded Theory is Not [Editorial]. *Academy of Management Journal*, *49*(4), 633-642. https://doi.org/10.5465/amj.2006.22083020

Sull, D., & Eisenhardt, K. M. (2015). *Simple Rules: How to Thrive in a Complex World*. Houghton Mifflin Harcourt.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning* (2nd ed.). MIT Press.

Tang, P. M., Koopman, J., McClean, S. T., Zhang, J. H., Li, C. H., Cremer, D., Lu, Y., & Ng, C. T. S. (2021). When conscientious employees meet intelligent machines: An integrative approach inspired by complementarity theory and role theory. *Academy of Management Journal*, *65*(3), 1019-1054. https://doi.org/10.5465/amj.2020.1516

Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. Alfred A. Knopf.

Tong, S., Jia, N., Luo, X., & Fang, Z. (2021). The Janus Face of Artificial Intelligence Feedback: Deployment Versus Disclosure Effects on Employee Performance. *Strategic Management Journal*, *42*(9), 1600–1631. https://doi.org/10.1002/smj.3322

Toorajipour, R., Sohrabpour, V., Nazarpour, A., Oghazi, P., & Fischl, M. (2021). Artificial intelligence in supply chain management: A systematic literature review [Article]. *Journal of Business Research*, *122*, 502-517. https://doi.org/10.1016/j.jbusres.2020.09.009

Torresen, J. (2018). A Review of Future and Ethical Perspectives of Robotics and AI [Review]. *Frontiers in Robotics and AI*, *4*(75). https://doi.org/10.3389/frobt.2017.00075

Townsend, D. M., Hunt, R. A., McMullen, J. S., & Sarasvathy, S. D. (2018). Uncertainty, Knowledge Problems, and Entrepreneurial Action. *Academy of Management Annals*, *12*(2), 659-687. https://doi.org/10.5465/annals.2016.0109

Townsend, D. M., Hunt, R. A., Rady, J., Manocha, P., & Jin, J. h. (2024). Are the Futures Computable? Knightian Uncertainty and Artificial Intelligence. *Academy of Management Review*, amr.2022.0237. https://doi.org/10.5465/amr.2022.0237

Tschang, F. T., & Almirall, E. (2021). Artificial intelligence as augmenting automation: Implications for employment. *Academy of Management Perspectives*, *35*(4), 642-659. https://doi.org/10.5465/amp.2019.0062

Tushman, M. L., & Anderson, P. (1986). Technological Discontinuities and Organizational Environments. *Administrative Science Quarterly*, *31*(3), 439-456. https://doi.org/10.2307/2392832

Tushman, M. L., & O'Reilly III, C. A. (1996). Ambidextrous Organizations: Managing Evolutionary and Revolutionary Change. *California Management Review*, *38*(4), 8-30. https://doi.org/10.2307/41165852

Ullman, T. D. (2023). Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks. *ArXiv*. Retrieved 15 Apr 2024, from https://arxiv.org/abs/2302.08399

Vanacker, T., Collewaert, V., & Zahra, S. A. (2017). Slack resources, firm performance, and the institutional context: Evidence from privately held European firms. *Strategic Management Journal*, *38*(6), 1305-1326. https://doi.org/10.1002/smj.2583

Vanneste, B., & Puranam, P. (2024). Artificial Intelligence, Trust, and Perceptions of Agency. *Academy of Management Review*, *In-Press*. https://doi.org/10.5465/amr.2022.0041

Verganti, R. (2008). Design, Meanings, and Radical Innovation: A Metamodel and a Research Agenda. *Journal of Product Innovation Management*, *25*(5), 436-456. https://doi.org/10.1111/j.1540-5885.2008.00313.x

Verganti, R., Vendraminelli, L., & Iansiti, M. (2020). Innovation and Design in the Age of Artificial Intelligence. *Journal of Product Innovation Management*, *37*(3), 212-227. https://doi.org/10.1111/jpim.12523

von Hippel, E. (1986). Lead Users: A Source of Novel Product Concepts. *Management Science*, *32*(7), 791-805. https://doi.org/10.1287/mnsc.32.7.791

von Hippel, E. (1994). 'Sticky Information' and the Locus of Problem Solving: Implications for Innovation. *Management Science*, *40*(4), 429-439. https://doi.org/10.1287/mnsc.40.4.429

von Hippel, E., Thomke, S., & Sonnack, M. (1999). Creating Breakthroughs at 3M. *Harvard Business Review*, *77*(5), 47-57.

Waardenburg, L., Huysman, M., & Sergeeva, A. V. (2022). In the Land of the Blind, the One-Eyed Man Is King: Knowledge Brokerage in the Age of Learning Algorithms. *Organization Science*, *33*(1), 59-82. https://doi.org/10.1287/orsc.2021.1544

Wagner-Tsukamoto, S. (2007). An Institutional Economic Reconstruction of Scientific Management: On the Lost Theoretical Logic of Taylorism. *Academy of Management Review*, *32*(1), 105-117. https://doi.org/10.5465/amr.2007.23463879

Watson, H. J. (2017). Preparing for the cognitive generation of decision support. *MIS Quarterly Executive*, *16*(3), 153-169.

Webb, T., Holyoak, K. J., & Lu, H. (2023). Emergent analogical reasoning in large language models. *Nature Human Behaviour*, *7*(9), 1526-1541. https://doi.org/10.1038/s41562-023-01659-w

West, D. M. (2018). *The Future of Work: Robots, AI, and Automation*. Brookings Institution Press.

Whang, O. (2023). Can a Machine Know That We Know What It Knows? *The New York Times*. Retrieved 05 Apr 2024, from https://www.nytimes.com/2023/03/27/science/ai-machine-learning-chatbots.html

Williams, J., Fiore, S. M., & Jentsch, F. (2022). Supporting Artificial Social Intelligence With Theory of Mind. *Frontiers in Artificial Intelligence*, *5*. https://doi.org/10.3389/frai.2022.750763

Wilson, H. J., & Daugherty, P. R. (2018). *Human + machine: Reimagining work in the age of AI*. Harvard Business Review Press.

Winfield, A. F. T. (2018). Experiments in Artificial Theory of Mind: From Safety to Story-Telling [Hypothesis and Theory]. *Frontiers in Robotics and AI*, *5*. https://doi.org/10.3389/frobt.2018.00075

Winter, S. G. (2012). Purpose and Progress in the Theory of Strategy: Comments on Gavetti. *Organization Science*, *23*(1), 288-297. https://doi.org/10.1287/orsc.1110.0696

Wright, S. A., & Schultz, A. E. (2018). The rising tide of artificial intelligence and business automation: Developing an ethical framework. *Business Horizons*, *61*(6), 823-832. https://doi.org/10.1016/j.bushor.2018.07.001

Wynne, B. E., & Dickson, G. W. (1975). Experienced managers' performance in experimental man-machine decision system simulation. *Academy of Management Journal*, *18*(1), 25-40. https://doi.org/10.5465/255622

Xu, H., Zhao, R., Zhu, L., Du, J., & He, Y. (2024). OpenToM: A Comprehensive Benchmark for Evaluating Theory-of-Mind Reasoning Capabilities of Large Language Models. *ArXiv*. Retrieved 08 Apr 2024, from https://arxiv.org/abs/2402.06044

Xu, Y., Verma, D., Sheridan, R. P., Liaw, A., Ma, J., Marshall, N. M., McIntosh, J., Sherer, E. C., Svetnik, V., & Johnston, J. M. (2020). Deep Dive into Machine Learning Models for Protein Engineering. *Journal of Chemical Information and Modeling*, *60*(6), 2773-2790. https://doi.org/10.1021/acs.jcim.0c00073

Zahedi, Z., Sengupta, S., & Kambhampati, S. (2024). 'Why Didn't You Allocate This Task to Them?' Negotiation-Aware Task Allocation and Contrastive Explanation Generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, *38*(9), 10243-10251. https://doi.org/10.1609/aaai.v38i9.28890

Zellweger, T. M., & Zenger, T. R. (2023). Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review*, *48*(3), 379-408. https://doi.org/10.5465/amr.2020.0503

Zhang, H., Ou, A. Y., Tsui, A. S., & Wang, H. (2017). CEO humility, narcissism and firm innovation: A paradox perspective on CEO traits. *The Leadership Quarterly*, *28*(5), 585-604. https://doi.org/10.1016/j.leaqua.2017.01.003

Zhou, E., & Lee, D. (2023). Generative AI, Human Creativity, and Art. *SSRN*. Retrieved 10 May 2024, from https://ssrn.com/abstract=4594824

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2020). Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey. *Working Paper Series - National Bureau of Economic Research*, *28290*. Retrieved 15 Jan 2024, from http://www.nber.org/papers/w28290

Zucker, L. G. (1987). Institutional Theories of Organization. *Annual Review of Sociology*, *13*(1), 443–464. https://doi.org/10.1146/annurev.so.13.080187.002303

Zuzul, T., & Tripsas, M. (2020). Start-up Inertia versus Flexibility: The Role of Founder Identity in a Nascent Industry. *Administrative Science Quarterly*, *65*(2), 395-433. https://doi.org/10.1177/0001839219843486

APPENDIX

Appendix to Chapter 1:
Methodology for the Systematic Literature Review

To obtain a holistic overview of the existing literature connecting strategic leadership and artificial intelligence, we searched the Web of Science for 9 unique terms: "strategic leadership", "leadership", and "management", each combined individually with each of the terms "artificial intelligence", "AI", and "machine learning". To ensure that the search results relate to modern developments, building on immense progress in software and hardware, i.e., the computing power enabling modern algorithms, we narrowed the search to the years 2000-2022. After filtering for categories (management, economics, information systems, artificial intelligence), citation topics (management, management science, economics, artificial intelligence), and publication types (articles, review articles), we obtained a total of 1,194 publications.

In a second step, we focused on the selection of leading journals in the field of management. For this, we combined four different and widely-known lists: The 50 journals used in the Financial Times research rankings, the top 50 SSCI Journal Impact Factor outlets in the fields of management, business, and management science, the ratings 4 and 4-star from the Association of Business Schools in the fields of general management, information management, marketing, operations and technology, management science, organization studies, and strategy, and finally the UT Dallas list. Removing the overlaps of journals appearing on multiple lists, this combination ultimately produced exactly 99 journals to be used in the analysis. Matching this list of journals to the collected publications resulted in 152 articles, i.e., articles using the relevant terminology, addressing the correct literature streams, and published in the selected journals.

We further complemented this top-down search approach with a manual, bottom-up approach, to ensure the inclusion of relevant papers that were not originally listed in the Web of Science search results. This bottom-up approach included three sources; i) a manual searching in leading academic- (i.e., the Academy of Management publications, Administrative Science Quarterly, Management Science, Organization Science, Strategic Management Journal) and practitioner outlets (California Management Review, Harvard Business Review, MIT Sloan Management Review, MISQ Quarterly Executive) for significantly wider search terms (e.g., management and innovation, creativity, digitization, technological change); ii) going through the bibliography of some of the previously identified papers on AI and Strategic Leadership, again focusing on these leading outlets and identifying other relevant papers in the conversation; iii) conversations with other academics on works they find most relevant to the conversation. This exercise led to the identification of an additional 44 papers, bringing the total list to 196.

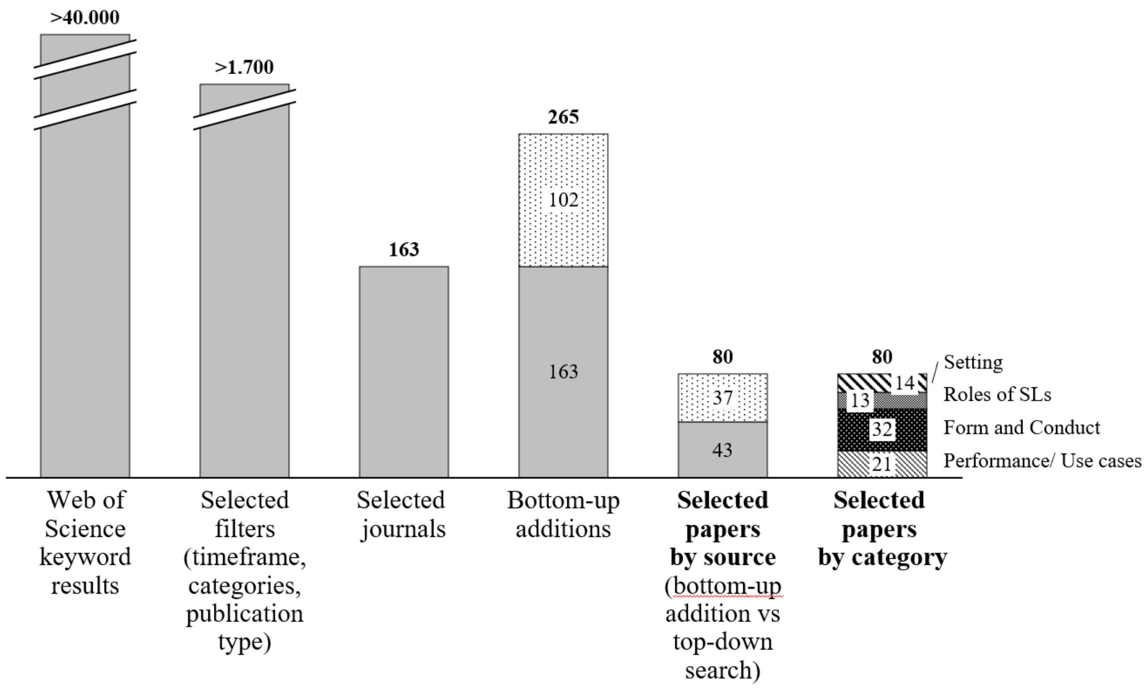

Figure 19: Number of papers identified through each step of the systematic literature review and by final category allocation

Out of these 196 papers, through studying abstracts and introductions and additionally skimming some full papers, we identified 68 papers squarely speaking to our topic of interest; the effects of AI on strategic leadership (see Figure 19). This list excludes many fantastic contributions to the discussion on AI, which lie outside of our focus. These are, for example, works on economic effects, which are cited in this chapter – but are not included in the list as they do not address leadership-specific aspects of AI. Figure 20 shows the distribution of papers across the years of publication, underlying the recent trend towards AI-specific literature in the field of management.

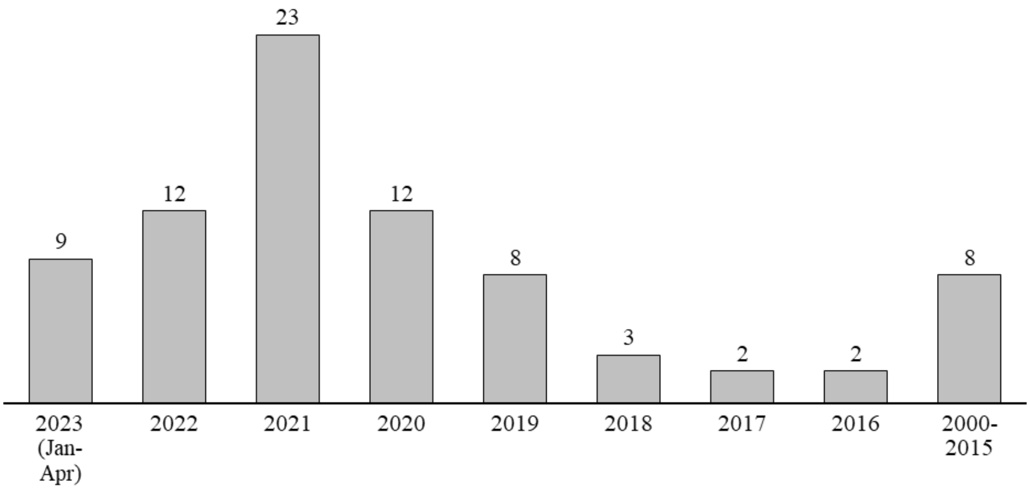

Figure 20: Number of papers included in the final list, by year of publication

We grouped the papers into four categories, matching the sections of the framework used in the analysis above. Of course, papers may bridge the boundaries of these categories, with papers addressing a change in the organizational setting, for example, also speaking to strategic aspects of the role of leaders. This made the division into categories a question of the highest degree of fit. The categories are: 1) Organizational setting, i.e., papers relating to the general management of AI, addressing questions such as how AI impacts organizations at large, or the capabilities of AI as a general purpose technology. 2) Roles of strategic leaders, i.e., papers relating to the practice of management, addressing the job functions and tasks of managers themselves, or the capability requirements for future leaders. 3) Organizational form and conduct, i.e., papers relating to shaping organizations around AI, including the design of processes within the organization, roles to be created, or the effective implementation of tools. 4) Organizational performance, i.e., identifying use cases suitable for AI tools, with papers presenting novel algorithms to solve specific problems.

Table 13: List of papers identified through the systematic literature review

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
		AI IN THE ORGANIZATIONAL SETTING					
1	Abada, I; Lambin, X	Artificial Intelligence: Can Seemingly Collusive Outcomes Be Avoided?	Management Science	2023	Empirical	Organizational Setting	Algorithms can reach collusive sales strategies when maximizing profits, requiring regulators to produce more socially desirable outcomes
2	Igna, I; Venturini, F	The Determinants Of AI Innovation Across European Firms	Research Policy	2023	Empirical	Organizational Setting	AI patent productivity positively correlates with a firms' previous innovation activity, indicating knowledge spillovers, especially for leading organizations
3	Krakowski, S; Luger, J; Raisch, S	Artificial Intelligence And The Changing Sources Of Competitive Advantage	Strategic Management Journal	2022	Empirical	Organizational Setting	AI can make humans' traditional competitive capabilities obsolete, while creating new sources of heterogeneity amongst competitors
4	Dixon, J; Hong, B; Wu, L	The Robot Revolution: Managerial And Employment Consequences For Firms	Management Science	2021	Empirical	Organizational Setting	Contrary to public opinion, robotics can increase employment, while reducing management positions and increasing spans of control
5	Gfrerer, A; Hutter, K; Füller, J; Ströhle, T	Ready Or Not: Managers’ And Employees’ Different Perceptions Of Digital Readiness	California Management Review	2021	Empirical	Organizational Setting	Perceptions of digital readiness and competences differ between management and employees
6	Haefner, N; Wincent, J; Parida, V; Gassmann, O	Artificial Intelligence And Innovation Management: A Review, Framework, And Research Agenda	Technological Forecasting And Social Change	2021	Conceptual	Organizational Setting	Framework showing the extent to which AI can replace humans and what is important in making the transformation
7	Raisch, S; Krakowski, S	Artificial Intelligence And Management: The Automation–Augmentation Paradox	Academy Of Management Review	2021	Review Essay	Organizational Setting	Augmentation and automation are interdependent and must be addressed jointly to avoid negative social outcomes
8	Shen, XR; Li, HS; Tolbert, PS	Converging Tides Lift All Boats: Consensus In Evaluation Criteria	Organization Science	2021	Empirical	Organizational Setting	Consensus in evaluation criteria of new technology increases investments to all firms in the new sector

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
		Boosts Investments In Firms In Nascent Technology Sectors					
9	Sturm, T; Gerlach, JP; Pumplun, L; Mesbah, N; Peters, F; Tauchert, C; Nan, N; Buxmannb, P	Coordinating Human And Machine Learning For Effective Organizational Learning	MIS Quarterly	2021	Empirical	Organizational Setting	Machine learning reduces need for explorative learning, but human intervention on ML has mixed effects
10	Tschang, FT; Almirall, E	Artificial Intelligence As Augmenting Automation: Implications For Employment	Academy Of Management Perspectives	2021	Conceptual	Organizational Setting	Automation and augmentation through AI will transform work, requiring critical conversations between business and society
11	Choudhury, P; Starr, E; Agarwal, R	Machine Learning And Human Capital Complementarities: Experimental Evidence On Bias Mitigation	Strategic Management Journal	2020	Conceptual	Organizational Setting	Predictions using ML tools can be worse than using prior technologies under influence of biased inputs
12	Hartmann, P; Henkel, J	The Rise Of Corporate Science In AI: Data As A Strategic Resource	Academy Of Management Discoveries	2020	Conceptual	Organizational Setting	The key drivers of AI development are large organizations, breaking the typical trend away from scientific research
13	Amabile, T	Guidepost: Creativity, Artificial Intelligence, And A World Of Surprises	Academy Of Management Discoveries	2019	Conceptual	Organizational Setting	AI becomes increasingly relevant in fields traditionally thought of as human domains, e.g., creativity or innovation
14	Kiron, D	What Managers Need To Know About Artificial Intelligence	MIT Sloan Management Review	2017	Conceptual	Organizational Setting	Managerial perspective on opportunities and risks emerging from AI
		AI IN THE PRACTICES OF STRATEGIC LEADERS					
15	Berg, J; Manav, R; Seamans, R	Capturing Value From Artificial Intelligence	Academy Of Management Discoveries	2023	Conceptual	Practices of Strategic Leaders	Understanding and developing complementary assets may be key to unlocking the potential of AI tools, specifically for large language models

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
16	Davenport, TH; Mittal, N	Stop Tinkering With AI It's Time To Go All In	Harvard Business Review	2023	Conceptual	Practices of Strategic Leaders	Creating a culture of data-driven decisions is one of the key challenges facing an organization, aside from tech-focused and structural requirements
17	Sun, JN; Zhang, DJ; Hu, HY; Van Mieghem, JA	Predicting Human Discretion To Adjust Algorithmic Prescription: A Large-Scale Field Experiment In Ware- House Operations	Management Science	2022	Empirical	Practices of Strategic Leaders	Human-centric algorithmics, taking into account human reactions to algorithmic prescriptions, may increase subsequent adherence and reduce errors
18	Wang, QC; Huang, Y; Jasin, S; Singh, PV	Algorithmic Transparency With Strategic Users	Management Science	2022	Methods	Practices of Strategic Leaders	Making algorithms transparent to the affected users may increase the predictive power of the algorithm, but also decrease overall performance.
19	Menz, M; Kunisch, S; Birkinshaw, J; Collis, DJ; Foss, NJ; Hoskisson, RE; Prescott, JE	Corporate Strategy And The Theory Of The Firm In The Digital Age	Journal Of Management Studies	2021	Conceptual	Leadership Practice	Practice of strategizing changes under the effect of digital technologies
20	Kruhse-Lehtonen, U; Hofmann, D	How To Define And Execute Your Data And AI Strategy	Harvard Data Science Review	2020	Conceptual	Leadership Practice	Executing an AI strategy necessitates ambition level, talent acquisition, and an adjusted operating mode
21	Lanzolla, G; Lorenz, A; Miron-Spektor, E; Schilling, M; Solinas, G; Tucci, CL.	Digital Transformation: What Is New If Anything? Emerging Patterns And Management Research	Academy Of Management Discoveries	2020	Conceptual	Leadership Practice	Effect of digital technologies on management practices is clear - but scale of shifts is unclear, requiring detailed analysis within the organization
22	Larson, L; DeChurch, LA	Leading Teams In The Digital Age: Four Perspectives On Technology And What They Mean For Leading Teams	Leadership Quarterly	2020	Conceptual	Leadership Practice	Implications of digital technologies on organizing and leadership in human-machine environments
23	Brock, JKU; von Wangenheimz, F	Demystifying AI: What Digital Transformation Leaders Can Teach	California Management Review	2019	Empirical	Leadership Practice	Guidance for leaders around understanding and effectively implementing AI fundamentals

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
		You About Realistic Artificial Intelligence					
24	Kiron, D; Schrage, M	Strategy For And With AI	MIT Sloan Management Review	2019	Conceptual	Leadership Practice	Strategy creation can differ between when working with AI versus working for AI
25	Merendino, A; Dibb, S; Meadows, M; Quinn, L; Wilson, D; Simkin, L; Canhoto, A,	Big Data, Big Decisions: The Impact Of Big Data On Board Level Decision-Making	Journal Of Business Research	2018	Empirical	Leadership Practice	Shortfall of capabilities on managerial level can impact team control and board-level cohesion
26	Kolbjørnsrud, V; Amico, R; Thomas, RJ	How Artificial Intelligence Will Redefine Management	Harvard Business Review	2016	Empirical	Leadership Practice	With automation coming for not only routine manual tasks, but also management, new must-have practices emerge
27	Wynne, Bayard E; Dickson, Gary W	Experienced Managers' Performance In Experimental Man-Machine Decision System Simulation	Academy Of Management Journal	1975	Empirical	Leadership Practice	Managers' personality traits affect their performance increases when working with machine systems
		AI IN ORGANIZATIONAL FORM AND CONDUCT					
28	Choi, WJ; Liu, QH; Shin, J	Predictive Analytics And Ship-Then-Shop Subscription	Management Science	2023	Empirical	Organizational Form and Conduct	With increasing prediction capabilities, firms can switch to ship-then-shop business models, providing consumers with a choice before paying, especially with larger search friction
29	Dodgson, M; Sheridan, A; Andrews, J; Phillips, N	Managing Technology-Enabled Innovation In A Professional Services Firm: A Cooperative Case Study	Academy Of Management Discoveries	2023	Empirical	Organizational Form and Conduct	Developing capilities for technology-enabled innovation is supported by leadership commitment and processes that take into account existing cultures and practices
30	Jia, N; Luo, X; Fang, Z; Liao, C	When And How Artificial Intelligence Augments Employee Creativity	Academy Of Management Journal	2023	Empirical	Organizational Form and Conduct	Higher-skilled employees may benefit from AI assistance as enhanced creativity leads to better performance

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
31	Lebovitz, S; Lifshitz-Assaf, H; Levina, N	To Engage Or Not To Engage With Al For Critical Judgments: How Professionals Deal With Opacity When Using Ai For Medical Diagnosis	Organization Science	2022	Empirical	Organizational Form and Conduct	Augmentation of knowledge work through AI is more likely for organizations allowing human experts to incorporate their own knowledge as well as AI projections
32	Giermindl, LM; Strich, F; Christ, O; Leicht-Deobald, U; Redzepi, A	The Dark Sides Of People Analytics: Reviewing The Perils For Organisations And Employees	European Journal Of Information Systems	2022	Conceptual	Organizational Form and Conduct	Increasingly analyzing human employees with AI tools can have negative effects on employees and organizations
33	Nguyen, TM; Malik, A	A Two-Wave Cross-Lagged Study On AI Service Quality: The Moderating Effects Of The Job Level And Job Role	British Journal Of Management	2022	Empirical	Organizational Form and Conduct	Satisfaction with AI services is affected by the adoption of AI tools in the workplace, moderated by job level
34	Waardenburg, L; Huysman, M; Sergeeva, AV	In The Land Of The Blind, The One-Eyed Man Is King: Knowledge Brokerage In The Age Of Learning Algorithms	Organization Science	2022	Empirical	Organizational Form and Conduct	Knowledge brokers translate algortihmic predictions, performing an essential function and receiving influence
35	Yang, JL; Chesbrough, H; Hurmelinna-Laukkanen, P	How To Appropriate Value From General-Purpose Technology By Applying Open Innovation	California Management Review	2022	Empirical	Organizational Form and Conduct	AI as a general purpose technology requires novel approaches for product go-to-market to appropriate value
36	Kim, S; Wang, Y; Boon, C	Sixty Years Of Research On Technology And Human Resource Management: Looking Back And Looking Forward	Human Resource Management	2021	Conceptual	Organizational Form and Conduct	Identification of patterns in research around HR and technology over 50 years
37	Kong, HY; Yuan, Y; Baruch, Y; Bu, NP; Jiang, XY; Wang, KP	Influences Of Artificial Intelligence (AI) Awareness On Career Competency And Job Burnout	International Journal Of Contemporary Hospitality Management	2021	Empirical	Organizational Form and Conduct	Use of AI tools in hospitality industry increases efficiency, but can have adverse effects, leading to burn-outs
38	La Torre, D; Colapinto, C;	Team Formation For Human-Artificial Intelligence Collaboration In The Workplace: A Goal	IEEE Transactions On Engineering Management	2021	Conceptual	Organizational Form and Conduct	Index of technology acceptance, technology self-efficacy, and source credibility to aid AI acceptance in teams

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
	Durosini, I; Triberti, S	Programming Model To Foster Organizational Change					
39	Li, JY; Li, MX; Wang, XC; Thatcher, JB	Strategic Directions For AI: The Role Of CIOs And Boards Of Directors	MIS Quarterly	2021	Empirical	Organizational Form and Conduct	Presence of a CIO positively influences AI orientation
40	Manesh, MF; Pellegrini, MM; Marzi, G; Dabic, M	Knowledge Management In The Fourth Industrial Revolution: Mapping The Literature And Scoping Future Avenues	IEEE Transactions On Engineering Management	2021	Conceptual	Organizational Form and Conduct	Analysis of knowledge management processes in the context of industry 4.0
41	Mikalef, P; Gupta, M	Artificial Intelligence Capability: Conceptualization, Measurement Calibration, And Empirical Study On Its Impact On Organizational Creativity And Firm Performance	Information & Management	2021	Empirical	Organizational Form and Conduct	AI capabilities support increase in organizational creativity and performance
42	Murray, A; Rhymer, J; Sirmon, DG	Humans And Technology: Forms Of Conjoined Agency In Organizations	Academy Of Management Review	2021	Conceptual	Organizational Form and Conduct	Four possible combinations of human and technological conjoined agency exist, each impacting routines differently
43	Pachidi, S; Berends, H; Faraj, S; Huysman, M	Make Way For The Algorithms: Symbolic Actions And Change In A Regime Of Knowing	Organization Science	2021	Empirical	Organizational Form and Conduct	Ignoring AI tools can lead employees to unwillingly support their implementation as successes are credited to the tool
44	Rahman, N; Daim, T; Basoglu, N	Exploring The Factors Influencing Big Data Technology Acceptance	IEEE Transactions On Engineering Management	2021	Empirical	Organizational Form and Conduct	Porposing factors to support the acceptance of AI tools witihn the organization
45	Smith, P; Beretta, M	The Gordian Knot Of Practicing Digital Transformation: Coping With Emergent Paradoxes In Ambidextrous Organizing Structures	Journal Of Product Innovation Management	2021	Empirical	Organizational Form and Conduct	In-depth analysis of one case organization's digital transformation, highlighting challenges and emerging tensions
46	Tang, PM; Koopman, J; McClean, ST.; Zhang, JH.; Li, CH;	When Conscientious Employees Meet Intelligent Machines: An Integrative Approach Inspired By Complementarity Theory And Role Theory	Academy Of Management Journal	2021	Empirical	Organizational Form and Conduct	More conscientious employees benefit less from the use of AI tools, raising questions on best usage in organizations

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
	de Cremer, D; Lu, Y; Ng, CTS						
47	Tong, S; Jia, N; Luo, X; Fang, Z	The Janus Face Of Artificial Intelligence Feedback: Deployment Versus Disclosure Effects On Employee Performance	Strategic Management Journal	2021	Empirical	Organizational Form and Conduct	AI-determined feedback increases quality, but leads to negative perceptions by employees
48	Balasubramanian, N; Ye, Y; Xu, M	Substituting Human Decision-Making With Machine Learning: Implications For Organizational Learning	Academy Of Management Review	2020	Conceptual	Organizational Form and Conduct	Danger of reducing diversity in routines, leading to myopia, through differences in human and machine learning
49	Glikson, E; Woolley, AW	Human Trust In Artificial Intelligence: Review Of Empirical Research	Academy Of Management Annals	2020	Conceptual	Organizational Form and Conduct	Employee trust in AI technology depends on circumstances such as representation, tangibility, and capabilities
50	Kellogg, KC; Valentine, MA; Christin, A	Algorithms At Work: The New Contested Terrain Of Control	Academy Of Management Annals	2020	Conceptual	Organizational Form and Conduct	Patterns of organizational control are re-shaped by AI, with potential dangers and pitfalls in implementation
51	Kronblad, C	How Digitalization Changes Our Understanding Of Professional Service Firms	Academy Of Management Discoveries	2020	Empirical	Organizational Form and Conduct	AI changes key aspects of Professional Service Firms, such as knowledge and capital intensity, leading to new practices
52	Schafheitle, S; Weibel, A; Ebert, I; Kasper, G; Schank, C; Leicht-Deobald, U	No Stone Left Unturned? Toward A Framework For The Impact Of Datafication Technologies On Organizational Control	Academy Of Management Discoveries	2020	Empirical	Organizational Form and Conduct	Framework to show how dataification technologies alter traditional control configurations
53	Wu, L; Hitt, L; Lou, BW	Data Analytics, Innovation, And Firm Productivity	Management Science	2020	Empirical	Organizational Form and Conduct	Data analytics capabilities more likely and more valuable in firms oriented around process improvement
54	Akhtar, P; Frynas, JG; Mellahi, K; Ullah, S	Big Data-Savvy Teams' Skills, Big Data-Driven Actions And Business Performance	British Journal Of Management	2019	Empirical	Organizational Form and Conduct	Big data savyness in teams leads to data-driven decision making and better performance

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
55	Shrestha, YR; Ben-Menahem, SM; von Krogh, G	Organizational Decision-Making Structures In The Age Of Artificial Intelligence	California Management Review	2019	Conceptual	Organizational Form and Conduct	Human and AI decision-making structures must be carefully combined to lead to optimal outcomes
56	Cockburn, I; Henderson, R; Stern, S	The Impact Of Artificial Intelligence On Innovation	Book: The Economics of AI	2018	Book chapter	Organizational Form and Conduct	Innovation and R&D processes reshaped through AI, leading to race for control of large datasets
57	Syam, N; Sharma, A	Waiting For A Sales Renaissance In The Fourth Industrial Revolution: Machine Learning And Artificial Intelligence In Sales Research And Practice	Industrial Marketing Management	2018	Conceptual	Organizational Form and Conduct	Impact of AI tools on sales practices and sales management
58	Singh, A; Hess, T	How Chief Digital Officers Promote The Digital Transformation Of Their Companies	MIS Quarterly Executive	2017	Empirical	Organizational Form and Conduct	Chief Digital Officers are becoming prevalent, but different possible roles necessitate careful evaluation of goals
59	Kiron, D; Kane, GC; Palmer, D; Phillips, AN; Buckley, N	Aligning The Organization For Its Digital Future	MIT Sloan Management Review	2016	Conceptual	Organizational Form and Conduct	People, processes, and culture must be aligned to achieve long-term digital success
		AI TOOLS IMPROVING ORGANIZATIONAL PERFORMANCE					
60	Unal, M; Park, YH	Fewer Clicks, More Purchases	Management Science	2023	Empirical	Tools Improving Performance	Using one-click buying functions on shopping websites persistently increases buying behavior for all customer types, likely through higher website engagement
61	Wang, W; Li, BB; Luo, XM; Wang, XY	Deep Reinforcement Learning For Sequential Targeting	Management Science	2023	Methods	Tools Improving Performance	Deep reinforcement learning may be used in marketing tools such as proce promotions to significantly increase long-term revenue by adjusting to consumers' price sensitivity
62	Hossain, MA; Agnihotri, R; Rushan, MRI;	Marketing Analytics Capability, Artificial Intelligence Adoption, And Firms? Competitive Advantage:	Industrial Marketing Management	2022	Empirical	Performance/ AI Use Cases	Marketing analytics capabilities as key driver for competitive advantage

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
	Rahman, MS; Sumi, SF	Evidence From The Manufacturing Industry					
63	Malik, A; Budhwar, P; Mohan, H; Srikanth, NR	Employee Experience -The Missing Link For Engaging Employees: Insights From An Mne'S AI-Based Hr Ecosystem	Human Resource Management	2022	Methods	Performance/ AI Use Cases	AI tools in HR management can enhance employee experience and engagement
64	Muhlroth, C; Grottke, M	Artificial Intelligence In Innovation: How To Spot Emerging Trends And Technologies	IEEE Transactions On Engineering Management	2022	Methods	Performance/ AI Use Cases	Identification of emerging trends and topics around new technological developments for organizations
65	Rusthollkarhu, S; Toukola, S; Aarikka-Stenroos, L; Mahlamaki, T	Managing B2B Customer Journeys In Digital Era: Four Management Activities With Artificial Intelligence-Empowered Tools	Industrial Marketing Management	2022	Empirical	Performance/ AI Use Cases	AI support in B2B customer journey management
66	Kim, SY; Upneja, A	Majority Voting Ensemble With A Decision Trees For Business Failure Prediction During Economic Downturns	Journal Of Innovation & Knowledge	2021	Methods	Performance/ AI Use Cases	Business failure prediction using decision tree algorithms
67	Mikalef, P; Conboy, K; Krogstie, J	Artificial Intelligence As An Enabler Of B2B Marketing: A Dynamic Capabilities Micro-Foundations Approach	Industrial Marketing Management	2021	Empirical	Performance/ AI Use Cases	Management of B2B marketing operations showing inter-relatedness of measures affecting value
68	Ozcan, S; Suloglu, M; Sakar, CO; Chatufale, S	Social Media Mining For Ideation: Identification Of Sustainable Solutions And Opinions	Technovation	2021	Methods	Performance/ AI Use Cases	Social media data mining to be utilised as a decision-making tool, detecting innovative ideas or solutions about a product
69	Liu, YZ; Qian, Y; Jiang, YC; Shang, J	Using Favorite Data To Analyze Asymmetric Competition: Machine Learning Models	European Journal Of Operational Research	2020	Methods	Performance/ AI Use Cases	Analysis of market competition, segmentation, and popularity of products within markets
70	Kumar, V; Rajan, B; Venkatesan, R; Lecinski, J	Understanding The Role Of Artificial Intelligence In Personalized Engagement Marketing	California Management Review	2019	Conceptual	Performance/ AI Use Cases	AI provides vast opportunities in personalized marketing to customers in an AI-driven environment

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
71	Metcalf, L; Askay, DA; Rosenberg, LB	Keeping Humans In The Loop: Pooling Knowledge Through Artificial Swarm Intelligence To Improve Business Decision Making	California Management Review	2019	Empirical	Performance/ AI Use Cases	ASI enables faster and improved group decision-making, harnessing diverse perspectives that often hinder convergence
72	Pandey, S; Pandey, SK	Applying Natural Language Processing Capabilities In Computerized Textual Analysis To Measure Organizational Culture	Organizational Research Methods	2019	Methods	Performance/ AI Use Cases	Text analysis and NLP tools used to identify organizational culture
73	Meyer, G; Adomavicius, G; Johnson, PE; Elidrisi, M; Rush, WA; Sperl-Hillen, JM; O'Connor, PJ	A Machine Learning Approach To Improving Dynamic Decision Making	Information Systems Research	2014	Methods	Performance/ AI Use Cases	Data mining classification to identify conditions for various AI-based decision-making strategies
74	Chen, ZY; Fan, ZP; Sun, MH	A Hierarchical Multiple Kernel Support Vector Machine For Customer Churn Prediction Using Longitudinal Behavioral Data	European Journal Of Operational Research	2012	Methods	Performance/ AI Use Cases	Customer churn prediction improves through use of novel algorithms
75	Das, SR; Chen, MY	Yahoo! For Amazon: Sentiment Extraction From Small Talk On The Web	Management Science	2007	Methods	Performance/ AI Use Cases	Improved sentiment analysis from text, e.g., for investor opinions on announcements
76	Cui, G; Wong, ML; Lui, HK	Machine Learning For Direct Marketing Response Models: Bayesian Networks With Evolutionary Programming	Management Science	2006	Methods	Performance/ AI Use Cases	Improved modeling consumer responses and other problems in marketing
77	Nissen, ME; Sengupta, K	Incorporating Software Agents Into Supply Chains: Experimental Investigation With A Procurement Task	MIS Quarterly	2006	Methods	Performance/ AI Use Cases	Analysis of boundary conditions for effective use of AI in supply chain management
78	Baesens, B; Verstraeten, G; Van den Poel, D; Egmont-Petersen, M;	Bayesian Network Classifiers For Identifying The Slope Of The Customer Lifecycle Of Long-Life Customers	European Journal Of Operational Research	2004	Methods	Performance/ AI Use Cases	Customer lifecycles analysis and purchase prediction through improved algorithms

#	Authors	Title	Source	Year	Type	Level of investigation	Summary
	Van Kenhove, P; Vanthienen, J						
79	Redmond, M; Baveja, A	A Data-Driven Software Tool For Enabling Cooperative Information Sharing Among Police Departments	European Journal Of Operational Research	2002	Methods	Performance/ AI Use Cases	Modelling of community profiles to aid police departments
80	Fowler, A	The Role Of AI-Based Technology In Support Of The Knowledge Management Value Activity Cycle	Journal Of Strategic Information Systems	2000	Conceptual	Performance/ AI Use Cases	Potential and limitations of AI tools in the knowledge management process

Appendix to Chapter 2: Data Structures and Data Analysis

Table 14: Preliminary data structure for case story – Blue

Source	Constructs *(exemplary quotes)*	Themes	Temporal brackets
	BLUE		
B2	*Build a central hub for analytics in the organization, from IT*	Individual goals as key drivers	Initiation
B2	*Central funding from the beginning*		
B2	*Department head and CIO were close and pulled it off together*		
B3	*CoE unit provide central infrastructure and PoC management in initial stages, but eventually became a bottleneck*	Arising issues	Exploration
B4	*Very many PoCs across the organization - but no tracking for results, no KPIs*		
B2	*The CIO left the company, and so the team remained, but the focus shifted to where it wasn't this high-pressure project anymore*	Emerging decentralization dynamics	
B5	*It became the wild west, everyone wanted to do something with data, everyone created azure or IBM accounts, the tech companies could sell their stuff separately to every department, we had robots everywhere, maximum hype*		
B4	*Cash no issue - if you had a fancy idea, you received the budget fairly easily*		
B4	*Became wasteful in terms of monetary resources and the capacities*	Need for action	Standardization
B5	*Some disillusionment: projects didn't scale as fast as imagined, producing little impact*		
B5	*Realizing that we've spent a lot of money, what did we effectively get in return? Trim things down and look for the more valuable use cases*		
B4	*New idea to create the digital strategy unit, in corporate strategy*	Push for coordination	
B1	*Aim to act outside of IT, centralistic, cross-functional, in touch with the entire organization*		
B5	*Decision-boards across multiple levels, aggregate upwards into the board level, using the same tools to document, report, and evaluate use cases*		
B5	*The new corporate strategy, driven by competitive environment, need to use the product opportunity, driven by R&D*	New communi-cation measures	
B5	*Started with large budgets for product development - but became a lighthouse effect, creating awareness*		
B4	*Clear measures to prioritize projects with the largest impact*		
B1	*Promote decentralized change while supporting centrally… Create willingness to invest in data science in business domains, much more efficient and sustainable*	Uncertainty in power and task allocation	
B4	*Definitely comparisons happening on board level - who has created targets, who has pilots. It makes them [executives] nervous to be behind*		
B2	*The hype is gone, we need to calculate and justify more financials*		
B1	*Change entire organization to become more software-driven, digital, and fast*	Emerging decentralization dynamics	Diffusion
B1	*Capabilities and network must be decentralized for AI, already reflected in the organization*		
B2	*People are expecting us [central team] to be gone soon*		
B3	*Topic was long driven by the IT, but now carried by businesses*		
B1	*Become the enabler and provide structure, rather than forcing people who might not even want to change*	New target roles for central teams	
B2	*Now teach data driven-thinking more than actual ideation*		
B1	*You have to convince people, it can't just be a mandate from the top or there will be resistance*		
B4	*The system has grown, [central units] do not have the reach to all projects anymore*		
B5	*Perspective changed; start with your own data instead of waiting for IT - only use them for implementation.*		
B5	*When we had little expertise, the central teams took all the best people and built their incubator*		
B1	*Takes years to promote this change, old systems and very diverse organization*	Power in silos as potential hurdle	
B5	*It's just less likely for someone to produce something completely useless - we all know the tools*		
B4	*Still some silos - it could help to have top managers say, 'I fully support this, I want to create transparency', but they're often trying to avoid conflicts*		

Table 15: Preliminary data structure for case story – Green

Source	Constructs - exemplary quotes	Themes	Temporal brackets
	GREEN		
G1	*[There was] fear of IT players entering the hardware space and taking over market share; reaction required on product-side*	External drivers for change	Initiation
G2	*They hype with deep learning, ImageNet, and those things in the consumer world, we saw those*		
G4	*while we are sending tables back and forth with suppliers. We have to find a better tool.*		
G2	*There are departments that have already gone quite far. They had use cases very early, started building up data analytics teams*	Independent development	Exploration
G1	*Urgency of response requires independent action by divisions, develop tailored solutions, supported by department management*		
G3	*We saw that they were already working in HR, in Controlling, in Finance - but very diverse approaches. That was interesting to see [in hindsight]*		
G3	*So much of this happened bottom-up, it was a grassroots-movement. We did have empowerment from the top, but there wasn't a corporate programme. It's really surprising, how bottom-up it all is.*		
G4	*It was always clear that digitization has to emerge in all different places in the organization, bottom-up. Individual use cases, individual ideas that act as small lighthouses within those areas*		
G1	*CTO initiated a project to assist the transition, central experts supported where requested*	Emerging need for action	Standardization
G2	*The question was; how do we get a measurable return on all this invest*		
G2	*How can we get AI more into the breath of the organization, more literacy*		
G2	*We have an inner source library and toolbox, with best practices to re-use for use case types, to approach scaling, become more efficient, become faster, and of course cheaper*	Coordination mechanisms	
G1	*Established a KPI system to track use cases, which is reported to CTO every 6 months*		
G3	*For our research teams I don't want to scale use cases, they should to be handed over [and controlled by central IT] after the prototype*		
G4	*Just from different sections using the same IT support, there is exchange on the technical side*		
G4	*Businesses have to buy in, at least have half the budget, and up to half from central funding*		
G2	*The hype had reached its peaked there. They were running down our doors,*	Communi-cation initiatives	
G4	*There are conversations in the different committees, with leadership from other departments, from corporate, regularly exchanging information, what could be interesting for us*		
G2	*[The businesses] do have to pay for use case support, they have to convince their leadership, so the top-down support is there*		
G1	*New CTO defines AI as key technology, driven as business targets for each unit*	Power dynamics between units	
G2	*There is definitely competition for talents - a central unit, with other talent, is more appealing. They don't want to be the single data scientist in some workshop, they want a network, develop. The central team gives them more research, more flexibility, less market demand*		
G3	*The biggest hurdle is still the silos between business units. One CEO doesn't want to talk to the other. One of them gets exchanged, and suddenly things work.*		
G1	*Good transparency exists on use cases, as far as possible in this large company.*	Long-term control requirements	
G5	*Will still need some centralization - in our case maybe more than we have right now. There are some red lines for tools that we can not cross and have to be aware of.*		
G6	*If everything would be decentralized, we would be re-inventing the wheel five times, every unit trying to solve the same problem, throwing money out the window*		
G1	*Can not be successful in the long run, if you only set things up centrally*	Long-term diffusion targets	Diffusion
G2	*The units will build up competency, and we have to ask why should a central department really exist*		
G2	*I hope that AI use will be standard, so that you won't even need us anymore. Maybe dedicated support within the verticals, but not centrally. That's the wrong focus*		
G5	*On 1=complete control to 10=complete chaos, we're a 6 right now within our global function. Tools are passed through global quality control, but regions implement things without our knowledge*	Business act independently	
G3	*There used to be these aliens, data scientists in the business, talking to the other aliens, data scientists in corporate. Today we have data scientists in the businesses themselves, with domain knowledge.*		
G4	*We need to be active in the daily business. No more step out of your role, go to a training- you lose speed. Do it internally, live at your work. Sharpen your axe while you are still busy cutting wood.*		
G3	*Some of our leadership has changed their thinking. We have these large visions, but it's hard to deduct any actions. So leadership has to endure a loss of control. As we say; Trust is the new control*	Central team losing purpose	
G2	*These question that require specific domain knowledge, that's not exactly the strength of a corporate research department. The value add is more second level support, if anything*		
G2	*We wanted to create synergies - but now the business focus is becoming more important. We are trying to facilitate an exchange between the silos, but that is now being questioned*		
G4	*If we were doing it centrally… I don't want to say they are lacking an understanding of the business, but at least a close relation to the business is missing*		

Table 16: Preliminary data structure for case story – Red

Source	Constructs - exemplary quotes	Themes	Temporal brackets
	RED		
R1	*It was directly initiated by the CIO, started in the old business intelligence unit*	External drivers for change	Initiation
R3	*First push into data analytics came from the central strategy and IT team, because the global, overarching trend was already clear. Then to set up a dedicated data and analytics team, for all departments and channels*		
R2	*Data [became more prominent] and the CIO really drove the initiative of creating a big data team - without a clear scope at the time*		
R1	*It could have been done more strategically. The impulse came from the top. But it happened fairly bottom-up.*	Emerging parallel structures	Exploration
R1	*We had different departments, the IT unit, the analytics team, and then partners out of Finance*		
R2	*It was more a living alongside each other than with each other.*		
R3	*Then created an overlapping set-up in the platform, where they used their own tech-stack, an own approach to the prioritization of use cases - and the roles and responsibilities between the teams were not clear*		
R2	*Not so much pull from business, more of an IT push into various areas*	Lack of coordination	
R2	*Often got stuck in the proof-of-concept phase, lacking involvement, not enough buy-in because of the IT push*		
R2	*There was lots of duplication of effort, lack of alignment, people thought it was cool to create own solutions*		
R2	*The team had data science, some engineering. But it didn't have data governance, ownership of the data sets, didn't have frontend. And there was also the separate unit in Finance, who supported analytics*		
R1	*It was a difficult time with IT [through the leadership changes], at that point the IT structure was slowing us down*	Need for action	Standardization
R2	*Priorities changed fairly quickly and so [with the change in leadership] there suddenly wasn't an owner for the topic any more*		
R2	*There was a clear strategic decision; Yes, we want to become a data-driven organization*	First attempt at more coordination	
R1	*Began picking out a few cherries to focus on. Still had the long-list in mind to think about what to do next*		
R2	*So the plan was to make it bigger, somehow, really end-to-end. With data-in, through the algorithms, to results in a dashboard - and move everyone we need to do that into Finance*		
R3	*What made it even more complicated: Engineering was with data analytics, but data science was with platform*		
R3	*We had a good start, but things changed and so organizationally it became a construct that was very difficult to understand*		
R1	*Covid was another shift - stronger focus on online sales, of course. So we looked how to support the platform*	Second attempt at more control	
R2	*The customer side became more important - but the data team was disconnected, because it's the other "kingdom". So the team had to move there*		
R2	*There still wasn't anything clearly defined in the corporate strategy or anything at the time*		
R1	*A stronger focus on data science and AI in the platform team with Covid, but now with the new strategy, we're back to investing more into the central unit. For every domain, the respective digital part is set to grow*	Third attempt/ success at establishing control	
R2	*The new strategy made clear that we are more B2C, the platform area will be stronger*		
R2	*For the first time, all areas of the organization are included, the strategy talks about digitization a lot.*		
R4	*If you take this new strategy seriously, it's a necessity to give data and analytics a new role*		
R4	*Focus on use cases that have a very clear buy-in from the top leadership. Only support our strategy and measured by the financial value. Not just thousand, but millions, or billions. That's the ambition*		
R4	*If someone is a data scientist and does data science activity, then they should be in our team in this stage. Nothing else in the rest of the organization*		
R2	*Depending on what the political weather is like, who is powerful at any time, they grab the topics that are hip and en-vogue. But then they can also go somewhere else again relatively quickly.*	Power dynamics shaping developments	
R1	*Why were platform and central IT always separate? That was just historically grown*		
R1	*We wanted the data science topic moved out of IT - it was always looked at too much though a cost-lens. something innovative like AI or data science needs more marketing.*		
R2	*Because of the vacancy on VP-level, you have these forces pulling the team apart. It's incredible how much politics is happening in these situations.*		
R2	*Again new leadership on board level who also wanted to get involved in data science*		
R3	*We had a CIO that was with the company for a year, and when they left, the IT unit moved again*		
R3	*Decisions are often made based on which division should have how much organizational power, it's politics*		
R3	*There were definitely also voices against [the importance of platform], but they just have more leverage*		
R1	*It's difficult, to be honest. Because yes, we do need cross-functional cooperation, to create synergies, but it does always slow us down. So we are trying to figure it out right now.*	Emerging issues with current set-up	Diffusion
R3	*If we don't start creating these processes and teams now, creating the readiness, then it might be the case that AI takes another step in it's technological evolution, and our processes won't be ready to deal with it*		
R2	*There is still a strong decentralization-pull; units saying, this is too central, we can't get to our own topics*		
R3	*The ambition is to be more vertical, with smaller market- and function teams*	More decentrali-zation as long-term goal	
R3	*It has to always be a product owner and product team, including data scientist, data engineer, dev-ops engineer - and all very close to the business partner*		
R4	*I hope that maybe in five years we'll be at the stage where we can say; ok, we might not need to centralize everything, we might move to business rather than a central organization*		
R3	*Product owners must not focus on the AI, but be responsible for the business problem. With reporting lines*		

Table 17: Description of organizational stages and exemplary quotes per case

	Model stages			
	Initiation	Exploration	Standardization	Diffusion
Duration (ca.)	year 0	year 1-3	year 3-6	year 6-9
Typical developments	Organization initiates work with AI, with external pressure or internal interest	AI interest emerges in various units, leading to first **decentralized** exploration of use cases with little central coordination or awareness	Central teams recognize potential and promote **centralized** monitoring and standardization of activities for synergies	Decision-making control handed back to units for **decentralized** development as interest scales beyond resource constraints of central teams
Blue	"Department head and CIO were close and pulled it off together" "Build a central hub for analytics in the organization, from IT"	"It became the wild west, everyone wanted to do something with data" "Cash was no issue - if you had a fancy idea, you received the budget" "Became wasteful in terms of monetary resources and the capacities"	"Idea to create the digital strategy unit" "Act outside of IT, centralistic, cross-functional, in touch with the entire organization" "Measures to prioritize projects with the largest impact"	"Capabilities and network must be decentralized for AI" "The system has grown, [central units] do not have the reach to all projects anymore" "People are expecting the central team to be gone soon"
Green	"Fear of IT players entering the hardware space and taking over market share required a reaction on the product side" "We saw the hype with deep learning and things in the consumer world, while we were sending hardcopy tables back and forth with suppliers. We have to find a better tool"	"Departments have already gone quite far. They had use cases very early, started building up data analytics teams" "Urgency of response requires independent action by divisions, develop tailored solutions" "So much of this happened bottom-up, it was a grassroots-movement"	"We established a KPI system to track use cases, which is reported to CTO every 6 months" "Teams are not allowed to scale their own use cases; they have to be handed over [and controlled by central IT] at some point" "There are conversations in the different committees, with leadership from other departments, from corporate, regularly exchanging information"	"Can not be successful in the long run, if you only set things up centrally" "Some of our leadership has changed their thinking […] endure a loss of control. As we say; Trust is the new control" "We wanted to create synergies - but now the business focus is becoming more important"
Red	"First push into data analytics came from the central strategy and IT team, because the global, overarching trend was already clear" "Set up a dedicated data and analytics team, for all departments and channels"	"The impulse came from the top. But [development] happened fairly bottom-up" "There was lots of duplication of effort, lack of alignment, people thought it was cool to create own solutions" "It was more a living alongside each other than with each other."	"There was a clear strategic decision; Yes, we want to become a data-driven organization" "So, the plan was to make it bigger, somehow, really end-to-end." "For the first time, all areas of the organization are included, the strategy talks about digitization a lot."	"Because yes, we do need cross-functional cooperation, to create synergies, but it does always slow us down." "There is still a strong decentralization-pull; units saying, this is too central, we can't get to our own topics"

<u>Table 18</u>: List of observed transitions and key descriptors

Tran-sition	Case	Giving up control	Taking control	Trigger	Successful?	Underlying power dynamic
1	**Blue**	Data Analytics (central)	Businesses	Organic	Short-term	No control mechanisms and budget independence for departments leads to diverse approaches
	Green	-	Businesses	Organic	Short-term	Ongoing independent product development in silos, informal alignment channels created
	Red	-	Platform / Data Analytics	Central plan	Yes	Active initiation by corporate strategy leading to resource ramp-up in separate teams
2	**Blue**	Businesses	Data strategy (central)	Inter-vention	Yes	Board reacts to lack of success and accountability
	Green	Businesses	Data strategy (central)	Inter-vention	Party	Board reacts to diverging paths, new CTO positions joint strategy as success
	Red	Platform	Finance/ Data (central)	Inter-vention	No	CFO takes control of IT, installs new CIO, data as prestige project
	Red	Finance/ Data (central)	Platform	Inter-vention	Short-term	Platform controls resources during Covid, CIO leaves company
	Red	Platform	Operations (central)	Central plan	Yes	New COO controls data team, pushes centralization and resource deliberation
3	**Blue**	Data strategy (central)	Businesses	Central plan	Yes	Planned redistribution of tasks between teams and ramp-up in businesses
	Green	Data strategy (central)	Businesses	Organic	Under radar start - *tbd*	Businesses begin ramp-up without central knowledge or involvement
	Red	Operations/ Data (central)	Businesses	Organic	*tbd*	Business ramp-up planned but not yet started – timing crucial

Appendix to Chapter 3: Data analysis and further regression models

Table 19: Correlation table for variables used in the final regression models (pairwise correlation, omitting NA values)

	Int. Depth	Int. Breadth	High-tech Industry	Orgsize	Orgage	Orgsuc	Inno-vation	IPrights	Dept. local	Dept. size	Dept. creat	Dept. digi	Dept. data	AI in other dept	Time AI use	Support	Support external	Dec. Distance	Board involv. [23]
Int. Depth	1.00	0.17	0.09	0.11	-0.13	0.32	0.02	0.26	0.12	0.28	0.51	0.51	0.50	0.17	0.29	0.30	0.26	-0.16	0.22
Int. Breadth	0.17	1.00	-0.10	-0.09	-0.17	0.07	0.06	0.07	-0.01	0.02	0.21	0.28	0.08	0.05	-0.03	-0.01	-0.02	0.19	-0.13
High-tech Industry	0.09	-0.10	1.00	0.10	-0.17	0.16	0.13	0.18	0.13	0.14	0.00	0.03	0.17	0.12	0.09	-0.02	0.05	-0.07	0.13
Orgsize	0.11	-0.09	0.10	1.00	0.30	0.12	-0.09	0.12	-0.07	0.44	-0.03	0.07	0.15	0.09	0.20	0.05	0.01	-0.20	0.31
Orgage	-0.13	-0.17	-0.17	0.30	1.00	-0.09	-0.14	-0.06	-0.08	0.03	-0.13	-0.11	-0.01	0.09	0.12	-0.27	-0.08	-0.01	-0.10
Orgsuc	0.32	0.07	0.16	0.12	-0.09	1.00	0.10	0.14	0.10	0.23	-0.02	0.22	0.24	0.10	0.20	0.17	0.08	-0.16	0.27
Innovation	0.02	0.06	0.13	-0.09	-0.14	0.10	1.00	0.13	-0.06	0.02	0.04	0.07	-0.01	-0.03	0.00	0.09	-0.05	-0.04	-0.02
IPrights	0.26	0.07	0.18	0.12	-0.06	0.14	0.13	1.00	0.06	0.12	0.03	0.10	0.40	0.03	0.10	-0.03	0.05	-0.13	0.15
Dept. local	0.12	-0.01	0.13	-0.07	-0.08	0.10	-0.06	0.06	1.00	0.14	-0.01	0.01	0.06	0.10	0.01	0.07	0.14	0.02	0.05
Dept. size	0.28	0.02	0.14	0.44	0.03	0.23	0.02	0.12	0.14	1.00	0.08	0.14	0.26	0.09	0.36	0.11	0.09	-0.21	0.41
Dept. creat	0.51	0.21	0.00	-0.03	-0.13	-0.02	0.04	0.03	-0.01	0.08	1.00	0.56	0.19	0.02	0.18	0.11	0.06	0.05	0.03
Dept. digi	0.51	0.28	0.03	0.07	-0.11	0.22	0.07	0.10	0.01	0.14	0.56	1.00	0.17	0.08	0.09	0.11	0.10	0.02	0.06
Dept. data	0.50	0.08	0.17	0.15	-0.01	0.24	-0.01	0.40	0.06	0.26	0.19	0.17	1.00	0.37	0.19	-0.02	0.20	-0.22	0.33
AI in other dept	0.17	0.05	0.12	0.09	0.09	0.10	-0.03	0.03	0.10	0.09	0.02	0.08	0.37	1.00	0.21	0.01	0.18	-0.07	0.22
Time AI use	0.29	-0.03	0.09	0.20	0.12	0.20	0.00	0.10	0.01	0.36	0.18	0.09	0.19	0.21	1.00	0.21	0.18	-0.05	0.22
Support	0.30	-0.01	-0.02	0.05	-0.27	0.17	0.09	-0.03	0.07	0.11	0.11	0.11	-0.02	0.01	0.21	1.00	0.34	-0.11	0.11
Support external	0.26	-0.02	0.05	0.01	-0.08	0.08	-0.05	0.05	0.14	0.09	0.06	0.10	0.20	0.18	0.18	0.34	1.00	-0.17	0.11
Dec. Distance	-0.16	0.19	-0.07	-0.20	-0.01	-0.16	-0.04	-0.13	0.02	-0.21	0.05	0.02	-0.22	-0.07	-0.05	-0.11	-0.17	1.00	-0.62
Board involv. [23]	0.22	-0.13	0.13	0.31	-0.10	0.27	-0.02	0.15	0.05	0.41	0.03	0.06	0.33	0.22	0.22	0.11	0.11	-0.62	1.00

[23] Decentralization Distance and Decentralization Board are highly correlated – these variables are not used in the same model, but instead as alternative measures for the decentralization of decision-making in models D and I respectively.

Table 20: Correlates of AI usage and depth of AI integration, all variations (model D)

(n=264 for all models) **Model D.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-2.99***	-2.99***	-3***	-3***	-3.01***	-3.02***	-3***	-3.01***	-3.01***	-2.99***
	(1.04)	(1.04)	(1.03)	(1.03)	(1.03)	(1.02)	(1.03)	(1.03)	(1.03)	(1.03)
High-tech industry [0/1]	1.31***	1.31***	1.31***	1.31***	1.31***	1.31***	1.31***	1.31***	1.31***	1.31***
	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)
Organization size [1-5]	0.1	0.1	0.11	0.11	0.11	0.1	0.11	0.11	0.11	0.1
	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)
Organization age [1-5]	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***	-0.4***
	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)
Org. success in previous years [1-5]	0.23*	0.23*	0.23*	0.23*	0.23*	0.23*	0.23*	0.23*	0.23*	0.23*
	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)
Innovation in products or processes [0/1]	2.01***	2.02***	2.01***	2.01***	2.01***	2.02***	2.01***	2.01***	2.01***	2.02***
	(0.37)	(0.37)	(0.37)	(0.37)	(0.37)	(0.37)	(0.37)	(0.37)	(0.37)	(0.36)
Importance of IP rights [1-4]	0.8***	0.8***	0.8***	0.8***	0.8***	0.8***	0.8***	0.8***	0.8***	0.8***
	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)
Regression on AI integration										
Intercept	-35.7***	-29.88***	-35.89***	-40.38***	-32.3**	-15.33	-34.92***	-38.55***	-33.29**	61.18
	(9.63)	(10.63)	(9.49)	(10.28)	(14.51)	(20.02)	(12.9)	(13.14)	(13.69)	(56.69)
High-tech industry [0/1]	-0.6	-0.83	-0.63	-0.57	-0.67	-0.92	-0.63	-0.55	-0.69	-1.11
	(2.25)	(2.25)	(2.2)	(2.2)	(2.2)	(2.18)	(2.2)	(2.21)	(2.2)	(2.24)
Organization size [1-5]	-0.12	-1.05	-0.16	-0.19	-0.83	-4.68 (3)	-0.15	-0.17	-0.56	-14.31
	(0.62)	(0.92)	(0.62)	(0.62)	(1.95)		(0.62)	(0.62)	(1.15)	(9.09)
Org. success in previous years [1-5]	2.27**	2.43**	2.11**	2.01**	2.07**	1.96**	2.13**	2.00**	2.07**	2.27**
	(0.96)	(0.97)	(0.95)	(0.95)	(0.95)	(0.98)	(0.98)	(0.97)	(0.97)	(1)
Innovation in products or processes [0/1]	-3.13	-3.35	-3.03	-2.77	-2.89	-2.63	-3.01	-2.74	-2.86	-2.46
	(3.84)	(4.01)	(3.86)	(3.78)	(3.81)	(3.98)	(3.86)	(3.77)	(3.81)	(4.25)
Importance of IP rights [1-4]	2.94	2.88	2.77	2.77	2.78	2.64	2.77	2.77	2.77	2.32
	(2.2)	(2.21)	(2.18)	(2.17)	(2.18)	(2.17)	(2.18)	(2.17)	(2.18)	(2.19)
Department located at HQ [0/1]	3.32	3.63	3.66	3.85	3.67	4.67	3.67	3.91	3.66	4.88*
	(2.9)	(2.92)	(2.9)	(2.95)	(2.89)	(2.99)	(2.89)	(2.94)	(2.88)	(2.87)
Department size [1-5]	1.1	1.17	0.79	0.85	0.78	0.89	0.8	0.85	0.78	0.73
	(0.88)	(0.88)	(0.9)	(0.9)	(0.91)	(0.9)	(0.9)	(0.9)	(0.9)	(0.87)
Department focus on creative tasks [% of tasks]	0.28***	0.27***	0.27***	0.27***	0.27***	0.25***	0.27***	0.27***	0.27***	0.25***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Department investment in digitization [% of budget]	0.24***	0.23***	0.25***	0.25***	0.25***	0.26***	0.25***	0.25***	0.25***	0.26***
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)

(n=264 for all models) **Model D.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Department availability of digital data [1-5]	9.02*** (1.66)	8.95*** (1.67)	9.09*** (1.66)	9.18*** (1.65)	9.1*** (1.66)	9.09*** (1.66)	9.05*** (1.69)	9.17*** (1.68)	9.07*** (1.7)	8.64*** (1.72)
AI usage in other departments [1-5]	-0.19 (1.03)	-0.17 (1.04)	-0.43 (1.01)	-0.45 (1.01)	-0.48 (1.01)	-0.48 (1.05)	-0.42 (1.01)	-0.45 (1.01)	-0.48 (1.01)	-0.47 (1.04)
Support structures used [0/1]	16.96*** (5.59)	16.99*** (5.54)	16.3*** (5.64)	16.86*** (5.61)	16.37*** (5.67)	17.03*** (5.53)	16.36*** (5.65)	16.9*** (5.64)	16.34*** (5.68)	15.32*** (5.36)
External support structures used [0/1]	2.21 (1.96)	2.13 (1.96)	1.96 (1.95)	1.87 (1.96)	2.04 (1.98)	1.45 (2)	1.96 (1.95)	1.84 (1.96)	2.07 (1.98)	1.73 (1.96)
Decentralization: Distance to board [0-x]	-0.52 (0.62)	-3.47 (2.31)	-0.59 (0.62)	1.37 (2.25)	-0.56 (0.62)	-10.59 (8.71)	-0.58 (0.62)	0.63 (1.34)	-0.57 (0.62)	-50.36** (22.77)
Time of AI usage in organization [1-5]			1.44 (1.01)	2.5* (1.39)	0.18 (3.39)	-2.9 (5.13)	0.7 (6.09)	0.96 (6.1)	1.26 (6.22)	-55.28 (36.26)
(Time of AI use)^2							0.12 (0.93)	0.27 (0.95)	-0.22 (1.15)	8.35 (5.44)
Org.Size x DecDistance		0.5 (0.37)				2.18 (1.5)				7.9** (3.92)
Org.Size x Time of use					0.24 (0.58)	1.03 (0.84)				7.95 (5.7)
Dec.Distance x Time of use				-0.62 (0.6)		1.92 (2.5)				31.17** (14.46)
Org.size x (Time of use)^2									0.05 (0.09)	-1.1 (0.84)
Dec.Distance x (Time of use)^2								-0.11 (0.08)		-4.85** (2.21)
Org.Size x Dec.Distance x Time of use						-0.47 (0.41)				-4.67** (2.36)
Org.Size x Dec.Distance x (Time of use)^2										0.69** (0.35)
Error terms										
sigma	13.01*** (0.6)	12.97*** (0.59)	12.95*** (0.6)	12.93*** (0.6)	12.94*** (0.6)	12.81*** (0.58)	12.95*** (0.6)	12.92*** (0.6)	12.94*** (0.6)	12.69*** (0.57)
rho	-0.2 (0.28)	-0.22 (0.3)	-0.19 (0.26)	-0.18 (0.25)	-0.18 (0.25)	-0.2 (0.27)	-0.19 (0.26)	-0.19 (0.25)	-0.18 (0.25)	-0.24 (0.31)
Adjusted R-squared	0.5544	0.5556	0.5551			0.5573	0.5552			0.5589

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

Table 21: Correlates of AI usage and AI integration breadth, all variations (Model B)

(n=264 for all models) **Model B.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)	-3.08*** (1.09)
High-tech industry [0/1]	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)	1.31*** (0.34)
Organization size [1-5]	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)
Organization age [1-5]	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)
Org. success in previous years [1-5]	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)	0.23* (0.13)
Innovation in products or processes [0/1]	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)	2.02*** (0.36)
Importance of IP rights [1-4]	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)
Regression on AI integration breadth										
Intercept	1** (0.44)	1.16** (0.47)	0.98** (0.43)	1.03** (0.46)	1.53** (0.68)	0.39 (0.79)	1.78*** (0.52)	1.82*** (0.52)	1.88*** (0.59)	-1.32 (2.25)
High-tech industry [0/1]	-0.24** (0.1)	-0.24*** (0.09)	-0.23** (0.09)	-0.23** (0.09)	-0.24** (0.09)	-0.24*** (0.09)	-0.24*** (0.09)	-0.24*** (0.09)	-0.24*** (0.09)	-0.23*** (0.09)
Organization size [1-5]	-0.03 (0.02)	-0.05 (0.04)	-0.03 (0.02)	-0.03 (0.02)	-0.13 (0.09)	0.07 (0.12)	-0.02 (0.02)	-0.02 (0.02)	-0.05 (0.04)	0.49 (0.41)
Org. success in previous years [1-5]	0.01 (0.03)	0.01 (0.03)	0.02 (0.03)	0.02 (0.03)	0.01 (0.03)	0.03 (0.03)	0.04 (0.03)	0.04 (0.03)	0.03 (0.03)	0.03 (0.03)
Innovation in products or processes [0/1]	-0.16 (0.32)	-0.17 (0.31)	-0.16 (0.31)	-0.17 (0.31)	-0.13 (0.32)	-0.15 (0.31)	-0.12 (0.3)	-0.13 (0.3)	-0.12 (0.3)	-0.14 (0.29)
Importance of IP rights [1-4]	-0.02 (0.08)	-0.02 (0.08)	-0.01 (0.08)	-0.01 (0.08)	-0.01 (0.08)	-0.01 (0.08)	-0.01 (0.08)	-0.01 (0.08)	-0.01 (0.08)	0.01 (0.08)
Department located at HQ [0/1]	-0.03 (0.12)	-0.02 (0.13)	-0.04 (0.12)	-0.05 (0.12)	-0.04 (0.11)	-0.05 (0.12)	-0.02 (0.11)	-0.03 (0.11)	-0.02 (0.11)	-0.02 (0.1)
Department size [1-5]	0.02 (0.03)	0.02 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.03 (0.03)	0.04 (0.03)
Department focus on creative tasks [% of tasks]	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002	0.002 (0.002
Department investment in digitization [% of budget]	0.005* (0.003)	0.005* (0.003)	0.004* (0.002)	0.004* (0.003)	0.004* (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)

(n=264 for all models) **Model B.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Department availability of digital data [1-5]	0.03 (0.06)	0.03 (0.06)	0.03 (0.06)	0.03 (0.06)	0.03 (0.06)	0.02 (0.06)	-0.01 (0.05)	-0.01 (0.05)	-0.01 (0.05)	0.01 (0.06)
AI usage in other departments [1-5]	0.03 (0.04)	0.03 (0.04)	0.04 (0.04)	0.04 (0.04)	0.03 (0.04)	0.03 (0.04)	0.04 (0.04)	0.04 (0.04)	0.04 (0.04)	0.03 (0.04)
Support structures used [0/1]	-0.08 (0.24)	-0.08 (0.24)	-0.05 (0.23)	-0.06 (0.22)	-0.04 (0.2)	-0.04 (0.18)	0.01 (0.19)	0 (0.19)	0.01 (0.19)	0.05 (0.15)
External support structures used [0/1]	-0.02 (0.06)	-0.03 (0.06)	-0.01 (0.06)	-0.01 (0.06)	0 (0.06)	0.02 (0.07)	0 (0.06)	0 (0.06)	0.01 (0.06)	0.01 (0.07)
Decentralization: Distance to board [0-x]	0.07*** (0.02)	-0.01 (0.09)	0.07*** (0.02)	0.05 (0.07)	0.07*** (0.02)	0.65** (0.32)	0.08*** (0.02)	0.06 (0.04)	0.08*** (0.02)	1.86** (0.85)
Time of AI usage in organization [1-5]			-0.05 (0.04)	-0.06 (0.06)	-0.24 (0.16)	0.17 (0.25)	-0.65*** (0.21)	-0.64*** (0.22)	-0.61*** (0.19)	1.34 (1.53)
(Time of AI use)^2							0.1*** (0.03)	0.09** (0.04)	0.08** (0.03)	-0.18 (0.24)
Org.Size x DecDistance		0.01 (0.01)				-0.1* (0.05)				-0.34** (0.15)
Org.Size x Time of use					0.03 (0.03)	-0.03 (0.05)				-0.32 (0.29)
Dec.Distance x Time of use				0.01 (0.02)		-0.2** (0.1)				-1.09** (0.56)
Org.size x (Time of use)^2									0.003 (0.004)	0.05 (0.05)
Dec.Distance x (Time of use)^2								0.001 (0.003)		0.15* (0.08)
Org.Size x Dec.Distance x Time of use						0.03** (0.02)				0.2** (0.1)
Org.Size x Dec.Distance x (Time of use)^2										-0.03* (0.01)
Heckman regression terms										
Inverse Mills Ratio (IMR)	0.18* (0.09)	0.18* (0.09)	0.17* (0.09)	0.18* (0.09)	0.17* (0.09)	0.16* (0.09)	0.17* (0.09)	0.17* (0.09)	0.16* (0.09)	0.15* (0.09)
Adjusted R-squared	0.1037	0.1028	0.1077			0.1301	0.1415			0.1583

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

Table 22: Correlates of AI usage and AI integration depth, measuring AI decentralization as the involvement of the board of management (model I)

(n=264 for all models) **Model I.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-2.99*** (1.04)	-2.98*** (1.05)	-3*** (1.04)	-3.01*** (1.04)	-3.01*** (1.03)	-3*** (1.04)	-3*** (1.04)	-3*** (1.04)	-3.01*** (1.03)	-2.98*** (1.06)
High-tech industry [0/1]	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)
Organization size [1-5]	0.1 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)
Organization age [1-5]	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)
Org. success in previous years [1-5]	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)
Innovation in products or processes [0/1]	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)
Importance of IP rights [1-4]	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)
Regression on AI integration depth										
Intercept	-38.84*** (10)	-39.17*** (10.17)	-39.37*** (9.88)	-38.45*** (9.88)	-34.83** (14.73)	-40.19*** (15.21)	-37.65*** (13.22)	-37.66*** (13.12)	-35.88** (13.97)	-59.56** (32.74)
High-tech industry [0/1]	-0.58 (2.28)	-0.68 (2.29)	-0.61 (2.22)	-0.48 (2.22)	-0.66 (2.22)	-0.63 (2.28)	-0.62 (2.23)	-0.49 (2.24)	-0.68 (2.22)	-0.55 (2.38)
Organization size [1-5]	-0.05 (0.62)	0.08 (0.69)	-0.08 (0.62)	-0.14 (0.62)	-0.95 (1.95)	0.5 (2.29)	-0.07 (0.62)	-0.13 (0.63)	-0.55 (1.15)	5.39 (6.46)
Org. success in previous years [1-5]	**2.35** (0.98)**	2.38** (0.98)	2.21** (0.97)	2.14** (0.96)	2.16** (0.97)	2.08** (0.95)	2.24** (1)	2.17** (0.99)	2.18** (0.99)	2.16** (0.98)
Innovation in products or processes [0/1]	-2.93 (3.82)	-3.15 (3.91)	-2.81 (3.83)	-2.46 (3.75)	-2.64 (3.76)	-2.5 (3.9)	-2.78 (3.83)	-2.47 (3.75)	-2.61 (3.76)	-2.12 (4.11)
Importance of IP rights [1-4]	3.02 (2.2)	3.02 (2.21)	2.87 (2.18)	2.88 (2.18)	2.89 (2.18)	2.85 (2.21)	2.87 (2.18)	2.89 (2.18)	2.88 (2.18)	2.71 (2.26)
Department located at HQ [0/1]	3.18 (2.89)	3.4 (2.94)	3.48 (2.9)	3.58 (2.93)	3.5 (2.88)	3.99 (2.95)	3.51 (2.88)	3.61 (2.92)	3.5 (2.88)	4.04 (2.92)
Department size [1-5]	1.2 (0.93)	1.24 (0.94)	0.91 (0.94)	0.98 (0.96)	0.91 (0.94)	1 (0.96)	0.92 (0.94)	0.99 (0.95)	0.91 (0.94)	1.03 (0.95)
Department focus on creative tasks [% of tasks]	**0.28*** (0.05)**	0.28*** (0.05)	0.27*** (0.05)	0.27*** (0.05)	0.26*** (0.05)	0.27*** (0.05)	0.27*** (0.05)	0.27*** (0.05)	0.26*** (0.05)	0.27*** (0.05)
Department investment in digitization [% of budget]	**0.23*** (0.06)**	0.23*** (0.06)	0.24*** (0.06)	0.24*** (0.06)	0.24*** (0.06)	0.23*** (0.06)	0.24*** (0.06)	0.24*** (0.06)	0.24*** (0.06)	0.23*** (0.06)

(n=264 for all models) **Model I.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Department availability of digital data [1-5]	**9.21*** **	9.23***	9.3***	9.26***	9.32***	9.28***	9.23***	9.22***	9.27***	8.91***
	(1.67)	(1.67)	(1.66)	(1.66)	(1.67)	(1.65)	(1.7)	(1.7)	(1.71)	(1.75)
AI usage in other departments [1-5]	-0.18	-0.26	-0.41	-0.47	-0.46	-0.82	-0.4	-0.47	-0.46	-0.7
	(1.06)	(1.09)	(1.03)	(1.03)	(1.03)	(1.11)	(1.03)	(1.04)	(1.03)	(1.1)
Support structures used [0/1]	**17.16*** **	17.08***	16.55***	16.94***	16.65***	16.79***	16.66***	16.97***	16.64***	16.4***
	(5.61)	(5.56)	(5.65)	(5.71)	(5.69)	(5.63)	(5.67)	(5.72)	(5.7)	(5.61)
External support structures used [0/1]	2.4	2.44	2.18	2.08	2.28	2.12	2.19	2.08	2.31	2.34
	(1.95)	(1.95)	(1.95)	(1.96)	(1.97)	(2)	(1.95)	(1.97)	(1.97)	(2)
Decentralization**: Board involvement** [0/1]	-0.58	4.8	-0.57	-7.54	-0.8	34.08	-0.62	-4.13	-0.82	166.41**
	(2.43)	(8.89)	(2.39)	(7.95)	(2.36)	(24.85)	(2.39)	(5)	(2.35)	(82.69)
Time of AI usage in organization [1-5]			1.36	0.92	-0.26	1.57	0.09	0.26	0.72	15.68
			(1.02)	(1.17)	(3.39)	(3.84)	(6.09)	(6.09)	(6.21)	(21.13)
(Time of AI use)^2							0.21	0.11***	-0.18	-2.22
							(0.93)	(0.94)	(1.14)	(3.18)
Org.Size x **Dec.Board**		-0.85				-7.2				-30.84**
		(1.39)				(4.43)				(13.38)
Org.Size x Time of use					0.3	-0.13***				-3.48
					(0.58)	(0.71)				(3.97)
Dec.Board x Time of use				2.23		-8.47				-102.18*
				(2.13)		(7.03)				(52.65)
Org.size x (Time of use)^2									0.05	0.52
									(0.09)	(0.58)
Dec.Board x (Time of use)^2								0.33		14.86*
								(0.33)		(7.85)
Org.Size x **Dec.Board** x Time of use						1.84				18.48**
						(1.2)				(8.23)
Org.Size x **Dec.Board** x (Time of use)^2										**-2.63*** **
										(1.21)
Error terms										
sigma	13.02***	13.02***	12.97***	12.94***	12.96***	12.88***	12.97***	12.94***	12.96***	12.83***
	(0.6)	(0.59)	(0.59)	(0.59)	(0.59)	(0.58)	(0.59)	(0.59)	(0.59)	(0.59)
rho	-0.21	-0.22	-0.19	-0.17	-0.18	-0.21	-0.19	-0.18	-0.19	-0.24
	(0.29)	(0.31)	(0.27)	(0.26)	(0.26)	(0.31)	(0.27)	(0.26)	(0.26)	(0.37)

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

Table 23: Correlates of AI usage and success of AI integration – focusing on Department AI readiness as correlate (model R)

(n=264 for all models) **Model R.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-3.01*** (1.03)	-3.01*** (1.03)	-3.02*** (1.03)	-3.02*** (1.03)	-3.03*** (1.02)	-3.03*** (1.02)	-3.02*** (1.03)	-3.02*** (1.03)	-3.03*** (1.02)	-3*** (1.02)
High-tech industry [0/1]	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)	1.31*** (0.33)
Organization size [1-5]	0.11 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)
Organization age [1-5]	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)	-0.4*** (0.13)
Org. success in previous years [1-5]	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)
Innovation in products or processes [0/1]	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.36)
Importance of IP rights [1-4]	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)	0.8*** (0.2)
Regression on AI implementation										
Intercept	-11.74 (9.15)	-6.12 (9.84)	-12.13 (9.07)	-16.24 (9.95)	-8.58 (14.54)	11.14 (19.94)	-8.98 (12.62)	-12.2 (12.95)	-7.28 (13.61)	89.45 (59.86)
High-tech industry [0/1]	0 (2.22)	-0.24 (2.22)	-0.06 (2.19)	-0.01 (2.19)	-0.11 (2.2)	-0.39 (2.18)	-0.09 (2.21)	-0.01 (2.21)	-0.15 (2.2)	-0.7 (2.24)
Organization size [1-5]	-0.16 (0.63)	-1.12 (0.91)	-0.19 (0.63)	-0.21 (0.63)	-0.85 (2)	-5.14* (2.98)	-0.17 (0.63)	-0.19 (0.63)	-0.59 (1.16)	-14.47 (9.46)
Org. success in previous years [1-5]	**2.41** (1.01)**	2.6** (1.02)	2.32** (1.02)	2.24** (1.01)	2.28** (1.01)	2.22** (1.05)	2.37** (1.03)	2.27** (1.02)	2.31** (1.03)	2.59** (1.03)
Innovation in products or processes [0/1]	-3.21 (3.35)	-3.42 (3.53)	-3.09 (3.3)	-2.83 (3.29)	-2.96 (3.26)	-2.63 (3.55)	-3.01 (3.31)	-2.74 (3.29)	-2.86 (3.27)	-2.35 (4.01)
Importance of IP rights [1-4]	**4.94** (2.1)**	4.9** (2.11)	4.82** (2.1)	4.84** (2.1)	4.83** (2.1)	4.7** (2.1)	4.78** (2.1)	4.8** (2.1)	4.78** (2.1)	4.04* (2.13)
Department located at HQ [0/1]	3.07 (3.12)	3.39 (3.15)	3.41 (3.12)	3.58 (3.16)	3.42 (3.1)	4.42 (3.16)	3.46 (3.1)	3.68 (3.14)	3.45 (3.09)	4.69 (3.01)
Department size [1-5]	1.38 (0.91)	1.45 (0.91)	1.07 (0.92)	1.13 (0.92)	1.06 (0.92)	1.17 (0.91)	1.09 (0.92)	1.14 (0.91)	1.07 (0.92)	0.97 (0.89)
Department AI readiness [0-100]	**0.55*** (0.05)**	0.54*** (0.05)	0.55*** (0.05)	0.55*** (0.05)	0.55*** (0.05)	0.55*** (0.05)	0.54*** (0.04)	0.55*** (0.05)	0.54*** (0.04)	0.53*** (0.05)
AI usage in other departments [1-5]	0.73 (0.93)	0.75 (0.93)	0.52 (0.91)	0.51 (0.92)	0.48 (0.93)	0.5 (0.95)	0.51 (0.92)	0.5 (0.92)	0.45 (0.93)	0.42 (0.95)

(n=264 for all models) Model R.	1	2	3	4	5	6	7	8	9	10
Support structures used [0/1]	**15.42*** (5.57)**	15.44*** (5.49)	14.65*** (5.57)	15.15*** (5.55)	14.71*** (5.6)	15.28*** (5.42)	14.9*** (5.59)	15.39*** (5.58)	14.87*** (5.63)	13.66*** (5.28)
External support structures used [0/1]	2.85 (1.99)	2.77 (1.98)	2.63 (1.98)	2.56 (1.99)	2.72 (2.02)	2.12 (2)	2.64 (1.98)	2.54 (1.99)	2.75 (2.02)	2.36 (1.96)
Decentralization: Distance to board [0-x]	-0.84 (0.62)	-3.89* (2.34)	-0.91 (0.63)	0.94 (2.21)	-0.89 (0.63)	-12.6 (8.59)	-0.88 (0.63)	0.27 (1.31)	-0.86 (0.62)	-56.61** (22.83)
Time of AI usage in organization [1-5]			1.46 (1.02)	2.45* (1.37)	0.21 (3.5)	-3.94 (5.24)	-1.07 (6.12)	-0.91 (6.17)	-0.53 (6.14)	-58.17 (38.39)
(Time of AI use)^2							0.41 (0.93)	0.56 (0.96)	0.07 (1.09)	8.67 (5.71)
Org.Size x DecDistance		0.51 (0.37)				2.44* (1.44)				8.98** (3.81)
Org.Size x Time of use					0.23 (0.6)	1.17 (0.85)				7.91 (5.98)
Dec.Distance x Time of use				-0.59 (0.59)		2.48 (2.52)				34.75** (14.68)
Org.size x (Time of use)^2									0.05 (0.09)	-1.07 (0.88)
Dec.Distance x (Time of use)^2								-0.11 (0.08)		-5.33** (2.24)
Org.Size x Dec.Distance x Time of use						-0.55 (0.4)				-5.31** (2.35)
Org.Size x Dec.Distance x (Time of use)^2										0.78** (0.35)
Error terms										
sigma	13.27*** (0.61)	13.23*** (0.59)	13.22*** (0.61)	13.19*** (0.61)	13.21*** (0.61)	13.06*** (0.59)	13.21*** (0.61)	13.18*** (0.61)	13.2*** (0.61)	12.91*** (0.58)
rho	-0.17 (0.22)	-0.18 (0.24)	-0.16 (0.21)	-0.16 (0.21)	-0.16 (0.21)	-0.19 (0.23)	-0.16 (0.21)	-0.16 (0.21)	-0.16 (0.21)	-0.23 (0.27)

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

Table 24: Correlates of AI usage and the percentage of key processes covered by AI (single dependent variable without PCF – model P)

(n=264 for all models) **Model P.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-2.91** (1.27)	-2.9** (1.27)	-2.94** (1.25)	-2.95** (1.22)	-2.96** (1.2)	-2.96** (1.17)	-2.93** (1.24)	-2.95** (1.2)	-2.94** (1.21)	-2.94** (1.17)
High-tech industry [0/1]	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)	1.3*** (0.33)
Organization size [1-5]	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)
Organization age [1-5]	-0.41*** (0.15)	-0.41*** (0.15)	-0.41*** (0.15)	-0.41*** (0.15)	-0.41** (0.15)	-0.4*** (0.14)	-0.41*** (0.15)	-0.41*** (0.15)	-0.41*** (0.15)	-0.41*** (0.14)
Org. success in previous years [1-5]	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)	0.22* (0.13)
Innovation in products or processes [0/1]	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)	2.01*** (0.37)
Importance of IP rights [1-4]	0.78*** (0.22)	0.78*** (0.22)	0.78*** (0.21)	0.78*** (0.21)	0.78*** (0.21)	0.79*** (0.21)	0.78*** (0.22)	0.78*** (0.21)	0.78*** (0.21)	0.78*** (0.21)
Regression on AI coverage of key processes										
Intercept	-26.69 (16.59)	-20.44 (17.61)	-27.09* (16.08)	-38.33** (16.04)	-18.89 (19.85)	-18.78 (25.88)	-13.94 (18.89)	-21.89 (18.17)	-12.15 (19.23)	32.98 (68.3)
High-tech industry [0/1]	-2.51 (3.12)	-2.79 (3.13)	-2.51 (3.02)	-2.33 (2.89)	-2.6 (2.92)	-2.78 (2.78)	-2.65 (3.1)	-2.39 (2.92)	-2.7 (3.02)	-2.9 (2.82)
Organization size [1-5]	0.62 (0.7)	-0.34 (1.09)	0.58 (0.7)	0.51 (0.69)	-0.97 (2.23)	-3.3 (3.5)	0.66 (0.7)	0.61 (0.69)	0.19 (1.3)	-7.61 (10.91)
Org. success in previous years [1-5]	1.22 (1.13)	1.39 (1.16)	1.06 (1.13)	0.82 (1.11)	0.96 (1.13)	0.89 (1.13)	1.31 (1.14)	1.06 (1.12)	1.24 (1.14)	1.24 (1.12)
Innovation in products or processes [0/1]	**-15.28** (7.53)**	-15.58** (7.52)	-15.11** (7.4)	-14.37* (7.53)	-14.75** (7.29)	-14.22* (7.45)	-14.89** (7.47)	-14.16* (7.52)	-14.7** (7.41)	-14.05* (7.62)
Importance of IP rights [1-4]	-1.24 (3.07)	-1.32 (3.1)	-1.39 (3.02)	-1.36 (2.9)	-1.36 (2.96)	-1.47 (2.85)	-1.43 (3.03)	-1.37 (2.88)	-1.42 (3)	-1.65 (2.88)
Department located at HQ [0/1]	2.48 (3.24)	2.8 (3.29)	2.85 (3.24)	3.32 (3.33)	2.87 (3.21)	4.17 (3.42)	3.01 (3.23)	3.51 (3.29)	3 (3.22)	4.42 (3.33)
Department size [1-5]	0.42 (0.94)	0.49 (0.94)	0.08 (0.99)	0.21 (0.96)	0.06 (0.99)	0.28 (0.96)	0.18 (0.98)	0.28 (0.96)	0.15 (0.99)	0.25 (0.94)
Department focus on creative tasks [% of tasks]	**0.38*** (0.07)**	0.37*** (0.07)	0.36*** (0.07)	0.36*** (0.07)	0.36*** (0.07)	0.34*** (0.07)	0.37*** (0.07)	0.36*** (0.07)	0.36*** (0.07)	0.35*** (0.07)
Department investment in digitization [% of budget]	**0.4*** (0.07)**	0.4*** (0.07)	0.41*** (0.07)	0.43*** (0.07)	0.41*** (0.07)	0.43*** (0.07)	0.4*** (0.07)	0.41*** (0.07)	0.4*** (0.07)	0.42*** (0.07)

(n=264 for all models) **Model P.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Department availability of digital data [1-5]	**6.46***** **(2.11)**	6.38*** (2.12)	6.54*** (2.1)	6.76*** (2.07)	6.56*** (2.11)	6.69*** (2.08)	6.04*** (2.09)	6.28*** (2.05)	6.06*** (2.09)	6.03*** (2.12)
AI usage in other departments [1-5]	-0.33 (1.29)	-0.32 (1.3)	-0.6 (1.27)	-0.67 (1.25)	-0.71 (1.26)	-0.78 (1.26)	-0.55 (1.25)	-0.62 (1.23)	-0.62 (1.25)	-0.7 (1.25)
Support structures used [0/1]	**18.22***** **(4.71)**	18.23*** (4.71)	17.53*** (4.71)	18.92*** (4.76)	17.7*** (4.69)	19.3*** (4.67)	18.33*** (4.7)	19.46*** (4.74)	18.3*** (4.71)	18.81*** (4.5)
External support structures used [0/1]	0.53 (2.28)	0.45 (2.31)	0.25 (2.28)	0.02 (2.26)	0.45 (2.3)	-0.17 (2.35)	0.37 (2.28)	0.1 (2.25)	0.49 (2.3)	-0.05 (2.32)
Decentralization: Distance to board [0-x]	-0.58 (0.75)	-3.64 (2.7)	-0.66 (0.75)	4.12* (2.42)	-0.6 (0.76)	-2.96 (9.57)	-0.55 (0.77)	1.93 (1.43)	-0.53 (0.77)	-25.47 (28.24)
Time of AI usage in organization [1-5]			1.54 (1.29)	4.13** (1.75)	-1.37 (3.92)	1.08 (6.19)	-8.18 (6.5)	-7.66 (6.42)	-7.54 (6.67)	-34.54 (43.49)
(Time of AI use)^2							1.57 (1.01)	1.87* (1.02)	1.19 (1.28)	5.63 (6.62)
Org.Size x DecDistance		0.52 (0.45)				1.41 (1.65)				4.28 (4.79)
Org.Size x Time of use					0.54 (0.68)	0.7 (1.01)				3.89 (6.97)
Dec.Distance x Time of use				-1.52** (0.67)		-0.95 (2.76)				15.06 (18.11)
Org.size x (Time of use)^2									0.05 (0.1)	-0.5 (1.05)
Dec.Distance x (Time of use)^2								-0.23** (0.09)		-2.55 (2.77)
Org.Size x Dec.Distance x Time of use						-0.16 (0.45)				-2.2 (2.95)
Org.Size x Dec.Distance x (Time of use)^2										0.33 (0.44)
Error terms										
sigma	14.86*** (0.76)	14.83*** (0.76)	14.8*** (0.74)	14.67*** (0.73)	14.78*** (0.74)	14.54*** (0.7)	14.74*** (0.75)	14.6*** (0.73)	14.73*** (0.75)	14.46*** (0.7)
rho	-0.18 (0.53)	-0.21 (0.55)	-0.16 (0.5)	-0.14 (0.45)	-0.15 (0.46)	-0.15 (0.42)	-0.18 (0.52)	-0.15 (0.45)	-0.17 (0.5)	-0.18 (0.44)

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*

Table 25: Correlates of AI usage and intensity of use, defined as no use/ testing/ rare use/ occasional/ regular use (single dependent variable without PCF – model Y)

(n=264 for all models) **Model Y.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Regression on AI Use										
Intercept	-3.1*** (1.01)	-3.1*** (1)	-3.1*** (1.01)	-3.1*** (1.01)	-3.1*** (1.01)	-3.11*** (1.01)	-3.1*** (1.01)	-3.1*** (1.01)	-3.1*** (1.01)	-3.1*** (1)
High-tech industry [0/1]	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)	1.32*** (0.33)
Organization size [1-5]	0.11 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.1 (0.09)
Organization age [1-5]	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)	-0.39*** (0.13)
Org. success in previous years [1-5]	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.24* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)	0.23* (0.12)
Innovation in products or processes [0/1]	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)	2.02*** (0.37)
Importance of IP rights [1-4]	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)	0.81*** (0.2)
Regression on AI usage intensity (no use/ testing/ rare use/ occasional use/ regular use)										
Intercept	-0.36 (0.4)	-0.2 (0.43)	-0.36 (0.4)	-0.29 (0.44)	-0.41 (0.58)	0.64 (0.8)	-0.72 (0.52)	-0.71 (0.53)	-0.68 (0.55)	3.76* (2.07)
High-tech industry [0/1]	0.04 (0.09)	0.04 (0.09)	0.04 (0.08)	0.04 (0.08)	0.04 (0.09)	0.04 (0.09)	0.05 (0.08)	0.04 (0.08)	0.04 (0.09)	0.03 (0.09)
Organization size [1-5]	-0.03 (0.03)	-0.05 (0.04)	-0.03 (0.03)	-0.03 (0.03)	-0.02 (0.07)	-0.19* (0.11)	-0.03 (0.03)	-0.03 (0.03)	-0.04 (0.04)	-0.65** (0.33)
Org. success in previous years [1-5]	**0.1*** (0.04)**	0.11*** (0.04)	0.1** (0.04)	0.1*** (0.04)	0.1** (0.04)	0.09** (0.04)	0.09** (0.04)	0.09** (0.04)	0.09** (0.04)	0.1** (0.04)
Innovation in products or processes [0/1]	**0.3* (0.18)**	0.3 (0.19)	0.31 (0.19)	0.3 (0.19)	0.3 (0.19)	0.31 (0.19)	0.3 (0.19)	0.29 (0.19)	0.3 (0.19)	0.31* (0.19)
Importance of IP rights [1-4]	0.23** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.22** (0.09)	0.2** (0.09)
Department located at HQ [0/1]	0.13 (0.11)	0.14 (0.11)	0.14 (0.11)	0.13 (0.11)	0.14 (0.11)	0.16 (0.12)	0.13 (0.11)	0.13 (0.11)	0.13 (0.11)	0.16 (0.11)
Department size [1-5]	0.06 (0.04)	0.06 (0.04)	0.05 (0.04)	0.05 (0.04)	0.05 (0.04)	0.05 (0.04)	0.04 (0.04)	0.04 (0.04)	0.04 (0.04)	0.04 (0.04)
Department focus on creative tasks [% of tasks]	**0.005*** (0.01)**	0.005*** (0.01)	0.005** (0.01)	0.005** (0.01)	0.005** (0.01)	0.004** (0.02)	0.005** (0.01)	0.005** (0.01)	0.005** (0.01)	0.004** (0.02)

(n=264 for all models) **Model Y.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Department investment in digitization [% of budget]	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
Department availability of digital data [1-5]	**0.35*** (0.07)**	0.35*** (0.07)	0.36*** (0.07)	0.35*** (0.07)	0.36*** (0.07)	0.35*** (0.07)	0.37*** (0.07)	0.37*** (0.07)	0.37*** (0.07)	0.34*** (0.07)
AI usage in other departments [1-5]	0 (0.04)	0 (0.04)	-0.01 (0.04)	-0.01 (0.04)	-0.01 (0.04)	0 (0.04)	-0.01 (0.04)	-0.01 (0.04)	-0.01 (0.04)	-0.01 (0.04)
Support structures used [0/1]	**0.47* (0.27)**	0.47* (0.27)	0.45* (0.27)	0.44 (0.27)	0.45* (0.27)	0.44* (0.26)	0.43 (0.27)	0.42 (0.27)	0.43 (0.27)	0.35 (0.27)
External support structures used [0/1]	0.12 (0.09)	0.12 (0.09)	0.11 (0.09)	0.12 (0.09)	0.11 (0.09)	0.1 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)	0.11 (0.09)
Decentralization: Distance to board [0-x]	-0.01 (0.03)	-0.1 (0.1)	-0.02 (0.03)	-0.05 (0.09)	-0.02 (0.03)	-0.57 (0.36)	-0.02 (0.03)	-0.02 (0.06)	-0.02 (0.03)	-2.33*** (0.75)
Time of AI usage in organization [1-5]			0.04 (0.04)	0.02 (0.05)	0.06 (0.13)	-0.22 (0.19)	0.31 (0.26)	0.31 (0.26)	0.33 (0.26)	-2.35* (1.34)
(Time of AI use)^2							-0.04 (0.04)	-0.04 (0.04)	-0.05 (0.04)	0.34* (0.2)
Org.Size x DecDistance		0.01 (0.02)				0.09 (0.06)				0.36*** (0.13)
Org.Size x Time of use					0 (0.02)	0.04 (0.03)				0.37* (0.21)
Dec.Distance x Time of use				0.01 (0.03)		0.15 (0.11)				1.46*** (0.48)
Org.size x (Time of use)^2									0 (0)	-0.05* (0.03)
Dec.Distance x (Time of use)^2								0 (0)		-0.22*** (0.07)
Org.Size x Dec.Distance x Time of use						-0.02 (0.02)				-0.22*** (0.08)
Org.Size x Dec.Distance x (Time of use)^2										0.03*** (0.01)
Error terms										
sigma	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.55*** (0.03)	0.54*** (0.03)
rho	-0.14 (0.15)	-0.15 (0.16)	-0.14 (0.15)	-0.14 (0.15)	-0.14 (0.15)	-0.15 (0.17)	-0.14 (0.15)	-0.14 (0.15)	-0.14 (0.15)	-0.17 (0.18)

*Robust standard errors in parentheses. Significance levels: * = 10%, **= 5%, ***= 1%*