NYENRODE BUSINESS UNIVERSITEIT

MBA Thesis

Creating value with agentic AI in marketing:

The view from marketing management

Daan Luttik

14 June 2026

Supervisor: S. Beninger

Nyenrode Business Universiteit

Straatweg 25, 3621 BG Breukelen

Foreword

Over 10 years ago I fell in love with what we then called AI. Back then I was a technical student writing an AI checkers computer that quickly became more capable at the game than I was. When I entered the field of marketing in 2018, AI became the hot new thing again; my first LinkedIn post ever showed me fumbling for an answer to the question “so what is AI?” Years later, after the releases of Dall-E and ChatGPT, the techniques that had interested me for so long came to the center of the public debate.

Additionally, when I graduated from my Master Data Science & Entrepreneurship I would never have guessed that only 5 years later I would start an Executive MBA. Nor would I have guessed that I would do an academic study again, now combining the key elements of my background in business, marketing, and AI.

I have greatly enjoyed my experience at Nyenrode, and this thesis is the crown to my entire academic journey. I want to thank the people who have brought me here: Techonomy and Peter Sprenger for giving me the opportunity to start the MBA. My team within the company for the great support they provided me whenever I had a busy period with the study.

Also, I want to thank all of my classmates for giving me such a great time during and outside of the module. And a very special thanks for my peer group, I love how we have supported each other. Another special thanks goes out to my family and friends, to those who I have disregarded because I was too busy, and those who supported me without expecting anything in return. I could not have made it through these busy years without having all of these special people around me to support me.

Lastly, I want to give a very special thanks to my supervisor Stefanie Beninger. During the module on marketing she taught, it quickly became clear to me that she was the perfect fit for this research. I want to thank you for your time and effort. And for the patience with all of the questions regarding your preferences, and the “are you sure” questions after you noted your preferences, I think it greatly improved my work.

Contents

Executive Summary

Generative AI has become central to marketing, which together with sales is projected to capture roughly a third of AI's economic value (Mayer et al., 2025). Yet most marketing AI still works as a question-and-answer tool (Brinker & Riemersma, 2025). A newer category, *agentic AI*, pursues goals autonomously across complex, multi-step tasks, promising to perform the work itself. How managers turn that promise into value is largely unexplored, and scholars have issued explicit calls for firm-level evidence (e.g. Kim, 2025). This study therefore asks: "How do marketing managers create value with agentic AI?" This study answers that question empirically using constructivist grounded theory (Charmaz, 2014). In doing so, this work provides one of the earliest empirical works on agentic AI within marketing.

This study proposes a model that combines four managerial behaviors: observing external conditions, navigating the organizational context, applying agentic AI, and obtaining value outcomes. The organizational context plays a particularly critical role, as it both enables and constrains the potential of agentic AI initiatives. We observe that marketing managers steer the marketing department, leverage technical resources, and deal with compliance policies.

Depending on the organizational context and managerial behaviors, identical use cases produced divergent outcomes across firms. A pattern of paradoxes runs throughout: the technology that creates a risk also supplies the means to govern it.

We offer advice to practitioners. Specifically, we suggest a set of managerial responses to the organizational context. Among other things, we advocate for a "laboratory" environment where marketing teams can safely experiment with agentic AI without violating compliance policies, and we note that marketing and technology managers should work together to negotiate the creation and allocation of resources.

As the field evolves quickly, many specific use cases and technical applications may date quickly; however, the higher-order principles regarding how marketing managers respond to this rapidly changing field are likely to be more durable.

1. Introduction

Generative artificial intelligence (GenAI) has rapidly become a central concern for marketing. Generative AI is a form of computer intelligence that can produce various forms of output, including text and images, based on patterns learned during training (Grewal et al., 2025). The use of GenAI presents significant business potential, and the field of marketing and sales may be the largest opportunity, representing about one-third of the estimated $4.4 trillion in annual productivity growth that might result from GenAI (Mayer et al., 2025).

In practice, marketers already apply GenAI across a range of tasks: drafting and localizing campaign copy, generating product and creative imagery, synthesizing customer insight from unstructured feedback, personalizing messages at scale, and powering conversational customer-service assistants (e.g. Grewal et al., 2025; Mogaji & Jain, 2024; Wahid et al., 2023). However, in marketing, AI still mainly serves as a question-and-answer machine, largely focusing on brainstorming and content generation rather than as a system that can really do the work or make a meaningful contribution to a vast number of everyday tasks (Brinker & Riemersma, 2025).

The field of AI evolves quickly, and recent innovations have led to agentic AI, a new category of GenAI which Acharya et al. (2025) define as "autonomous AI systems that undertake to finish a set of complex tasks that span over long periods of time without human supervision" (p. 18913). These systems are capable of acting in a goal-oriented manner within a changing environment (Acharya et al., 2025) and might therefore be better suited to the execution of the marketer's normal daily tasks.

These technological developments might enable marketers to leverage agentic AI for this large number of daily tasks not yet affected by AI, which in turn might unlock the large productivity gains mentioned by Mayer et al. (2025). However, the use of agentic AI in marketing remains largely empirically unexplored, and many researchers are calling for research on agentic AI within the marketing field (Kim, 2025) or for research related to the broader field of GenAI in marketing (Jain et al., 2024; Mogaji & Jain, 2024). Kim (2025) notes that the transformational shift that agentic AI brings requires a serious response from the research community. This study therefore explores the question "How do marketing managers create value with agentic AI?" The question's exploratory nature reflects the field, which is still emerging and in which academic understanding and practical guidance remain underdeveloped. To guide the research, the following sub-questions are used:

1. How do marketing managers identify and translate agentic AI opportunities?
2. How do marketing managers adopt agentic AI within their organizations?
3. How do marketing managers manage the benefits, sacrifices, and risks that relate to the utilization of agentic AI?

This study deliberately focuses on agentic AI because of its ability to act autonomously. Within his call for research, Kim (2025) notes the ability of agentic AI to integrate data, make decisions, and execute campaigns autonomously, which seems to align with the untapped potential that Brinker and Riemersma (2025) describe. This is where understanding is thinnest.

Existing accounts of digital transformation offer a useful starting point regarding how to adopt these new techniques, since the rollout of AI can be read as a new wave of such transformations like those described by Holmström (2022), but they were not written for technologies that act with this degree of autonomy and evolve this quickly, and so cannot be assumed to transfer wholesale.

This study makes both theoretical and practical contributions. Theoretically, it offers one of the first grounded empirical accounts of how marketing managers create value with agentic AI. Its central contribution is a model of managerial value creation that links observing the external conditions, navigating the organizational context, applying agentic AI, and obtaining the resulting value outcomes. Furthermore, this study describes where beneficial or counterproductive conditions and managerial behaviors align with those observed in earlier work within the related fields of GenAI and digital transformation.

Practically, the study offers managers insights into use cases they might apply, the conditions that mediate value creation with agentic AI, and managerial behaviors they can apply to shape the conditions for successful implementation, adoption, and use of agentic AI. Furthermore, this work describes potentially beneficial conditions and behaviors that marketing managers are typically unaware of. Additionally, Appendix B provides a practical checklist that managers might apply to audit their organizational readiness and shape their agentic AI initiatives.

The remainder of this document is structured as follows: First, an overview of the current state of the literature is provided, including calls for research in this domain. Then an overview of the methods is provided and the findings are presented. Afterward, the findings are compared to literature, practical implications are derived, limitations are discussed, and directions for further research are provided. Lastly, a conclusion is provided containing the main points.

2. Literature review

2.1. The historical development of AI

The term Artificial Intelligence (AI) has been around since the 1950s to describe "intelligent machines" (Xu et al., 2021). It is often associated with systems that perform tasks that are typically associated with human cognition (Xu et al., 2021). AI can be broadly categorized into two categories:

1. Expert systems: Where rules and heuristics written by humans define the output of the system (Kastner & Hong, 1984). An example of this can be seen in early games where playing against the computer was noted as playing against the AI.
2. Machine learning: Where data is used to train a model (Shinde & Shah, 2018).

Both expert systems and machine-learning methods have been applied in marketing decision-making for decades. Earlier expert systems were developed to support managerial decisions through data, models, rules, and analytical tools (Little, 1979; Mitchell et al., 1991). In contemporary marketing practice, the rule-based logic of expert systems is especially visible in marketing automation systems, where marketers define parameters, heuristics, lead-scoring rules, behavioral triggers, and customer segments to automate or personalize marketing actions such as email campaigns, content delivery, and sales-lead qualification (Guercini, 2022; 2023; Heimbach et al., 2015; Järvinen & Taiminen, 2016).

Machine learning is also widely used in marketing, particularly for customer segmentation, forecasting, recommender systems, consumer-behavior prediction, and content analysis (Duarte et al., 2022). Specific applications of machine learning include online A/B testing, where machine learning can be used to optimize the balance between exploration and exploitation, and conversion optimization (Kohavi et al., 2009), product recommendation systems in e-commerce (Schafer et al., 2001), and customer lifetime value modeling, where predicted future customer value can be compared with acquisition and retention costs to support marketing investment decisions (Bauer & Jannach, 2021; Gupta et al., 2004).

Deep learning is a specific form of machine learning that uses neural networks consisting of multiple layers (LeCun et al., 2015). This approach became more prevalent after AlexNet, which significantly improved the ability to recognize the primary object within an image compared to the previous state of the art (Krizhevsky et al., 2012). Probably the best-known early applications of deep learning in marketing revolve around chatbots. After word2vec (Mikolov et al., 2013), a technique that provides numerical representations of words and enabled superior natural language processing by machines, it became possible to match user questions to explicitly collected question-and-answer sets. This gave rise to early chatbots, which could answer well-documented questions (Dutta et al., 2021).

2.1.1. The rise of the LLM

However, the real breakthrough came with the paper "Attention is All You Need" (Vaswani et al., 2017), which introduced the Transformer model, a new technique that made AI much better at analyzing text. The later term GPT is a technically descriptive name; it stands for generative pretrained transformer (Radford et al., 2018) and refers to this technique. Training these systems became costly (Cottier et al., 2024) and pretrained models became more prevalent. One of the best-known models from before ChatGPT was BERT in 2018 (Devlin et al., 2019); this model was often applied to categorize text or apply sentiment analysis (Li et al., 2025). Within the realm of marketing, this was highly valuable to analyze customer support conversations or product reviews. At the same time, OpenAI worked toward increasingly larger models, starting with GPT-1 (Radford et al., 2018) and becoming widely known with GPT3 (Brown et al., 2020) and ChatGPT (OpenAI, 2022).

Models have become generally applicable and can therefore be trained before deciding on a specific application, hence the term "pretrained". The pretrained nature drastically changes the economic model (Burkhardt & Rieder, 2024). Rather than requiring vast amounts of data and having to recoup the cost of training, organizations can now directly start experimenting with the application of models (Burkhardt & Rieder, 2024). Using pretrained models shifts the costs from being upfront to usage-based costs, making AI much more applicable for, amongst others, small-scale use cases and quick experimentation (Burkhardt & Rieder, 2024).

Furthermore, while BERT focuses on classification (Devlin et al., 2019), GPT focuses on generation; the ability to generate text makes the model more widely applicable since it is no longer scoped to one specific task (Burkhardt & Rieder, 2024).

2.1.2. The impact of ChatGPT

ChatGPT was the introduction of GenAI to the broader public as a simple question-and-answer machine (OpenAI, 2022). ChatGPT can generate text based on the user's input; this, in addition to the earlier release of stable diffusion and other imaging models, made the public keenly aware of the potential for AI to answer questions and create content (Wahid et al., 2023). Marketers started using the image generation to make better briefings for the design department and used the text systems to write better copy, be a brainstorming partner for the creative process, and much more, all the while this is done more effectively (Elgheit, 2025; Grewal et al., 2025; Wahid et al., 2023).

Vaid et al. (2025) note that GenAI is also starting to be used for marketing tasks like segmentation and targeting. However, Vaid et al. (2025) do not describe what was needed to deploy these use cases, nor what business outcomes resulted from the addition of GenAI to these processes.

2.1.3. The transition to agentic AI

Just before the release of ChatGPT, the concept of Retrieval Augmented Generation (RAG) was introduced (Lewis et al., 2020). RAG supplements a model's input with relevant text retrieved from external knowledge sources (Lewis et al., 2020). This development was immediately relevant for businesses to put additional information into the context of the model. The additional information loaded through RAG can be anything, including business- or project-related content or publicly available information (Gao et al., 2023). RAG powers the AI system with additional knowledge that is typically directly related to the posed prompt (Gao et al., 2023; Lewis et al., 2020); this allows GenAI systems to answer questions using the information from your organization.

The ability to act was a vital next step in the development of agentic AI. A key conceptual step was the ReAct framework, introduced by Yao et al. (2023), which showed how LLMs could combine reasoning (Re) and acting (Act) in an iterative process, leveraging the results of an action to consider what a next step might be. The ability to consider actions and affect external systems enabled AI agents to act on behalf of the user and affect business applications (Hughes et al., 2025).

After the ReAct framework, many innovations, standards, and implementations that have built on this concept have improved the quality and availability of AI agents. This includes technical means like planning and external memory, early implementations like function calling in OpenAI (OpenAI, 2023), and the release of standards like the Model Context Protocol (Anthropic, 2024) and Skills standards (Anthropic, 2025). These standards have created a situation where specific agentic capabilities, like task execution, have become common in consumer-grade AI software. For instance, the skills standard is now supported in tools like Claude and ChatGPT (ChatGPT, 2026; Claude, 2025).

The ability to interact with other systems, often referred to as *tool calling*, is the foundation for agentic AI. It allows for AI to perform a much broader set of activities (Hughes et al., 2025). Nowadays, AI is capable of setting goals, planning to achieve these goals, and performing the steps of the plan (Hughes et al., 2025). This includes retrying and changing the plan when hurdles or unexpected input are found.

This allows AI to not only be usable for content generation and for brainstorming but also to interact with traditional software (Hughes et al., 2025). This, in turn, provides the basis for a world where marketers can use AI with the systems they already use, like marketing automation platforms, customer relationship management systems, and the like (Hughes et al., 2025). Hughes et al. (2025) also note applications of agentic AI that are not specific to marketing, including improving and automating business processes and developing actionable insights through data analysis.

2.2. On Agentic AI

While the technological progress towards agentic AI is highly concrete, the definition of agentic AI varies. Bandi et al. (2025) identified 12 aspects where agentic AI might differ from traditional AI; its primary function was defined as "goal-oriented autonomous decision-making with multi-step workflow execution." Moralles et al. (2026) note that literature converges on four primary characteristics: autonomy, goal-directedness, adaptivity, and proactivity. In one sentence, Moralles et al. (2026) define it as follows: "An agentic system not only acts independently but does so in pursuit of objectives it selects, constructs, or interprets."

In this research, the definition from Moralles et al. (2026) is followed, which has been simplified in communications with interviewees to two key points: 1. The ability to act and 2. The ability to take steps in pursuit of an objective.

While Thapa et al. (2026) describe the potential of agentic AI for marketing campaign execution from a theoretical perspective, little academic research has examined how agentic AI is used within marketing in practice. There are, however, many calls for research on this topic. Mogaji and Jain (2024) and Jain et al. (2024) call for further research on how generative AI is changing consumer behavior, marketing practice, and policy. Kim (2025) makes a more specific call for research into advertising in the age of agentic AI. These calls are consistent with the gap identified in this thesis: academic literature has begun to address generative AI in marketing but has not yet developed sufficient empirical understanding of how marketing managers create value with agentic systems beyond content generation.

Recent work on organizational readiness for agentic AI shows the importance of the organizational context to the adoption of agentic AI. Srivastava (2026) provides a conceptual, non-empirical overview of organizational readiness criteria for agentic AI, which include strategic readiness, governance & accountability readiness, workforce & leadership readiness, data architecture, and boundary readiness.

Furthermore, Schmidt et al. (2026) describe the concept of *process debt* as a key concept for leveraging agentic AI, they note that if you do not know how a process works (e.g. decision criteria are hidden, data is not readily available), it becomes difficult to delegate that process to an agent. However, the research by Schmidt et al. (2026) did not focus on the context of marketing, did not describe how the manager can shape these readiness criteria, and Srivastava (2026) called for research to describe the relationship between the organizational readiness for agentic AI and the realized value of the agentic systems.

Lastly, Hasselwander et al. (2026) describe a key direction in which agentic AI is evolving. Increasingly more functionality and services are exposed through agentic AI systems, this results in a broadly applicable agentic AI system that they call the AI super assistant (Hasselwander et al., 2026). Systems like these might upend how one might execute "digital tasks" (Hasselwander et al., 2026) and in doing so will likely also affect how companies interact with consumers and how marketers perform their work.

2.3. Value theory and value creation with agentic AI

Because the central research question asks how marketing managers create value, the thesis requires a clear but flexible understanding of value. Marketing theory treats value as a multidimensional construct. Woodside et al. (2008) emphasize that customer value is composed of multiple components; one of the formulae that they provide defines value as the "relative sum of total weighted benefits perceived – relative total costs perceived." This indicates that value is (1) a balance between benefits and costs, and (2) is based on perception.

Furthermore, Woodside et al. (2008) show that benefits and costs can take multiple forms, like "product benefits" or "emotional benefits". Almquist et al. (2016) expand this perspective and provide an overview with 30 different forms of value grouped in the categories, one of which is the “functional” category, which includes values like “saves time” and “reduces effort”. Other categories address different kinds of needs; for example, the “emotional” category contains “entertainment” and the “life changing” category contains “self-actualization” (Almquist et al., 2016).

An example of how this context can apply to AI is provided by Doshi and Hauser (2024); while not analyzing agentic AI, they implicitly note benefits and costs of using GenAI for creative work. They note that when using GenAI for creative work, individual creativity goes up (benefit); however, it also reduces collective novelty (cost).

2.4. Learnings from earlier waves of digital transformation

Fernandez-Vidal et al. (2022) describe the process of digital transformation, noting that it generates strategic and operational changes in companies resulting from the opportunities and threats of digital technologies. Since agentic AI is a novel digital technology, it makes sense to analyze it through this lens.

A consistent lesson is the *digital paradox*: organizations frequently fail to extract value from digital technologies despite heavy investment, because technology value is mediated by complementary assets: the skills, data, processes, structure, and leadership that surround the tool (Ancillai et al., 2023). Enholm et al. (2022) and Holmström (2022) explicitly link AI adoption to digital transformation. Holmström (2022) frames AI readiness as exactly this based on similar elements: the capabilities, data, governance, and learning practices an organization must build before AI produces value. Enholm et al. (2022) show a link between AI Capabilities and competitive value, highlighting technological, organizational, and environmental enablers and inhibitors and noting efficiency as a key first-order effect, among others. Note that both the work by Enholm et al. (2022) and Holmström (2022) was done before GenAI became prevalent and well before agentic AI came into play.

Fernandez-Vidal et al. (2022) describe digital transformation from the managerial perspective. While focusing mainly on technical roles, Fernandez-Vidal et al. (2022) highlight the importance of management during digital transformations, noting its role in driving business change, managing technical talent, and prioritizing learning.

Gebauer et al. (2020) describe a concerning pattern which they call the digitalization paradox; this pattern notes that while companies often invest significantly in digital transformation, they often do not yield the expected returns. Gebauer et al. (2020) note key traps in the process of digitization, like ignoring customer needs, applying incomplete accounting, and failing to build trust among partners.

2.5. Synthesis: sensitizing concepts for the empirical study

The literature reviewed above suggests that agentic AI in marketing should be examined within the context that it operates in, rather than as a standalone technology. The technical literature explains why agentic AI is different: it can combine LLM reasoning, memory, tool use, and interaction to pursue goals across workflows. The marketing literature explains why the context matters: marketing already contains many AI-relevant activities, but value depends on the connection between technology, customer knowledge, operations, and strategic decision-making. The digital transformation literature explains why adoption is difficult: organizations must go beyond technical capabilities, considering contextual factors and shaping organizational factors before digital technologies reliably produce value. The value literature explains why the outcome is multidimensional: value may be operational, customer-related, strategic, or co-created, and it may be accompanied by value destruction.

3. Methodology

3.1. Research approach

Given the exploratory nature of the research question and the emerging character of agentic AI in marketing, constructivist grounded theory, a form of qualitative research, was applied (Charmaz, 2014). Constructivist grounded theory enables researchers to develop theory grounded in participants’ experiences, meanings, and actions, while explicitly recognizing that analysis is shaped by context, interaction, and the researcher’s interpretive role (Charmaz, 2014). In turn, these insights were collected and converted into theory.

Grounded theory has the researcher iterate and switch between the collection and analysis of qualitative data; this leads to surprises and sparks of ideas (Charmaz, 2006). These surprises and sparks of ideas are especially relevant in a new and evolving field like agentic AI, where the researchers do not yet have a robust understanding of where the most relevant insights might be found.

Blumer (1986) notion of sensitizing concepts was used to spark ideas to pursue, as they can provide a relevant start to the research (Charmaz, 2014). A key sensitizing concept for this research is the concept of value, particularly as described in the pyramid of value by Almquist et al. (2016) and the various definitions of value as noted by Woodside et al. (2008).

3.2. Data collection

The process for data collection within constructivist grounded theory is flexible but well-defined, combining interviews with memo writing and coding (Charmaz, 2014).

The research employed in-depth interviews with two key stakeholder groups: 1. marketing managers, who bring the domain expertise and practical experience, and 2. AI experts who can add additional expert knowledge about the technical possibilities and limitations of agentic AI solutions. This dual perspective was chosen to get additional insights on both the technological possibilities and the current and planned use within marketing. The interaction between these perspectives can draw out insights that would otherwise be missed and further inform theory and practice.

In line with constructivist grounded theory, a semi-structured interview was used (Charmaz, 2014). To guide initial interviews, a semi-structured interview guide and protocol were created (see Appendix A).

The process of iteratively collecting and analyzing data within constructivist grounded theory starts with open, intensive interviews (Charmaz, 2014). Afterward, the data is analyzed through the coding process. In this process, statements from the interviews are categorized (Charmaz, 2014); this forms the basis for further analysis. Coding, in its three forms (initial, focused, and theoretical), forms the link between the interviews and the resulting theory (Charmaz, 2014). An overview of the coding process is found in Table 4, and a list of definitions for the theoretical codes is provided in Table 3.

Throughout the process, memos were written to capture the researcher's thoughts (Charmaz, 2014). Once tentative categories were established using the coding process, theoretical sampling was applied, i.e., instead of sampling for population representativeness, specific data is sought to elaborate and refine the properties of the emerging categories (Charmaz, 2014). After reaching theoretical saturation, the point where new data no longer yields new insights (Charmaz, 2014), the memos and categories were utilized to write the concrete arguments that form the basis for the new theoretical contribution (Charmaz, 2014).

3.3. Sample and sample size

In pursuit of the research question, top managers and experts in the field were interviewed. For the first group, the marketing managers, a slightly wider scope than just those with the formal title was used to identify relevant candidates: the focus is on those who have a clear managerial role in the marketing team, e.g. a CMO or marketing manager, but also those who are leading AI initiatives within the marketing team, i.e. they might not be the direct supervisor, but they are leading the change process or implementation. The second group consists of AI experts; they were selected based on knowledge of AI and expertise in implementing and rolling out AI systems. This group contains both technical leaders and technical consultants.

To find sufficient participants, a combination of theoretical sampling (Charmaz, 2014) and snowball sampling (Parker et al., 2019) was used to identify interview participants. The sampling started with the researcher's personal network, relying mainly on LinkedIn; candidates were selected when their roles matched the selection criteria. Since the sampling universe was already narrow and access was limited, no further selection criteria were applied to potential participants. However, key properties[1] about the participant's role and organization were noted to facilitate analysis of insights and potential limitations derived from the sample distribution. Data gathering continued until new data no longer yielded new theoretical insights. During data gathering, a circular process of administering was applied including interviews, writing memos, and coding data into relevant categories that informed subsequent interviews (Charmaz, 2014).

Table 1 contains an anonymized overview of the candidates, including their background, role, and type of company. Due to the network of the researcher and the research institution, most of the interviewees worked mainly in the Netherlands (11 out of 17 total).

3.4. Trustworthiness

This study is assessed for trustworthiness using the criteria from Lincoln and Guba (1985): credibility, transferability, dependability, and confirmability. To address this topic, strategies adapted from Shenton (2004) for interview-based projects of limited duration were followed. The details of this assessment are included in table form (see Table 2).

The researcher's academic background in AI and professional experience in marketing and AI provided domain fluency that aided rapport and probing, but also introduced a risk of bias towards overemphasizing the importance of technical aspects.

[1]These include, for the company, industry, year of founding, number of employees, and for the employee, the role and age. All numerical values are bucketed to allow for sufficient anonymity of the participants.

Table 1: Overview of participants.

Interviewee number	**Background**			**Primary role**			**Company type**		
	Technical Data	Technical AI	Marketing	Marketing manager	Internal AI expert	External AI advisor	Marketing agency/ consultant	Internal marketing	Other advisory role
1		x				x	x		
2			x	x				x	
3	x		x	x		x		x	
4			x			x		x	
5	x				x			x	
6			x	x				x	
7						x			x
8			x		x			x	
9			x		x		x		
10			x			x	x		
11			x			x	x		
12			x			x	x		
13	x	x				x	x		
14			x			x	x		
15			x	x				x	
16	x	x	x		x			x	
17			x			x	x		

Table 2: Trustworthiness assessment

Criteria of trustworthiness	**How addressed**
Credibility concerns the congruency between the findings and the phenomenon studied (Merriam, 1998; Shenton, 2004).	The study applied constructivist grounded theory (Charmaz, 2014) with iterative interviewing, memo writing, and coding. Credibility was strengthened through triangulation across marketing practitioners and AI experts; iterative initial and focused coding (Table 4); theoretical sampling; retention of negative and paradoxical accounts; the researcher had prior familiarity with the field and applied reflexivity (Section 3.4); and regular thesis-supervisor debriefing on transcripts, memos, and emerging categories. Member checking was not conducted (Charmaz, 2014). Theoretical saturation was assessed during data collection: later interviews chiefly densified existing categories rather than introducing new top-level concepts, with seventeen interviews completed (Table 1).
Transferability requires sufficient contextual detail for readers to judge whether findings may apply elsewhere, not statistical generalization to a wider population (Lincoln & Guba, 1985; Shenton, 2004).	In line with Shenton (2004), information regarding the sampling method, participant types, and data collection methods are provided (Sections 3; Table 1; Appendix A). Thick description in Chapter 4 and the limitations in Section 5.3 allow readers to judge whether findings may apply in comparable settings.
Dependability concerns whether the inquiry process is consistent, auditable, and reported in enough depth that a future researcher could follow the same steps (Shenton, 2004).	Interviews were transcribed; transcriptions, memos, and emerging models were reviewed in regular debriefing meetings with the thesis supervisor (Shenton, 2004). Anonymized data were submitted to the research institution for examination. The research design, field procedures, and reflective appraisal are documented in Chapter 3 and Appendix A (Shenton, 2004). Triangulation across marketing managers and AI experts provided overlapping methods. The code structure is reported in Table 4 and supporting quotes are provided.
Confirmability requires that findings emerge from participant accounts rather than researcher predisposition (Lincoln & Guba, 1985; Shenton, 2004).	In line with Shenton (2004), analytic memos recorded reasoning separately from transcript text; preliminary interpretations that were not supported by later data were revised or discarded. Findings claims are anchored in participant quotations (Chapter 4). Reflexive documentation of the researcher's marketing and AI background and sampling implications appears in Section 3.2.

Table 3: Definitions of the theoretical categories

Theoretical category	**Definition**
Observing external conditions	The managerial behavior of monitoring and interpreting the changing context outside of the organization.
Navigating the organizational context	How managers steer or move around conditions within the organization.
Applying agentic AI	Using or implementing agentic AI systems within the organization, across key marketing use cases.
Obtaining value outcomes	The value obtained from the agentic AI initiative.

Table 4: Overview of code development

Open coding	Focused coding	Theoretical coding
AI's improvement & potential; Observing increasing AI adoption; vendor roadmaps & offerings	Observing AI progression	Observing external conditions
Competitor pressure; Changing consumer behavior; The rise of agents of consumers	Observing Market pressure	
Resistance to change; Limited AI literacy; Lacking systems thinking; Analysis paralysis; Absent innovation culture; Strategic direction & vision; Senior leadership buy-in & backing; Educating & training; Experimenting with AI; Bringing people along; Providing clarity; Providing leadership backing; Championing AI	Steering the marketing department	Navigating the organizational context
Lacking data & infrastructure; Lacking technical talent; Leveraging data availability; Leveraging the right tooling; Leveraging external experts & vendors	Leveraging technical resources	
Restrictive compliance & legal gatekeeping; Politics & working in silos; Navigating restrictive governance	Dealing with compliance	
Generating insights	Generating insights	Applying agentic AI
Creating & validating content & campaigns	Creating & validating content & campaigns	
Using generic agents; Automating processes; Leveraging tool calling; Personalizing agents	Utilizing generic agents for personal work	
Deploying customer-facing agents	Employing customer-facing agents	
Gaining efficiency & speed; Gaining scale; Extending the personal skillset; Improving quality of output	Obtaining Benefits	Obtaining value outcomes
Incurring financial costs; Displacing jobs	Obtaining Sacrifices	
Risking hallucination; Risking security & privacy violations; Risking brand degradation	Obtaining Risks	

4. Findings

The findings indicate a wide range of applications and use cases of agentic AI in marketing. Additionally, a clear pattern of managerial behaviors emerged that enables and stimulates the adoption of agentic AI and mediates value creation from agentic use cases. Furthermore, notable impact is observed stemming from the organizational context (internal and external) on the adoption of agentic AI and the value creation from agentic AI.

Figure 1 presents the resulting process model. It shows how the studied marketing managers observed an impetus for change in their external environment; navigated the organizational conditions that enable or constrain the implementation and use of agentic AI within the organization; applied agentic AI through specific use cases; and obtained value outcomes.

4.1 Observing external conditions

The process of creating value using agentic AI starts with the manager observing changing conditions that signal a pressure or an opportunity to change. *Observing* refers to the managerial process of monitoring and interpreting the organization's changing context. With regard to agentic AI, two key categories of external conditions that sparked change have been identified: the progression of AI itself and market pressures.

This observation process typically leads both marketing managers and AI experts to recognize the significant impact of agentic AI. As interviewee 7 noted, "We're living in an era of change, and we're living in a change of era simultaneously." Interviewee 6 began incorporating AI across all parts of her organization, noting, "So I'm of course responsible for commercial, but … this is the future for our company!"

4.1.1 Observing AI progression

The first external condition that the marketing manager might observe is AI progression. This concept describes how AI is evolving in the marketplace; improving in quality, gaining new capabilities, and becoming more widely adopted. Managers observe this pattern; how they respond is developed in section 4.2. Marketing managers and AI experts consistently describe agentic AI as fast-moving. Interviewee 14 noted "These developments are going so fast, they really go in a hockey stick. … What was spectacular within Gemini is taken over all at once by Claude the day after."

Interviewee 8, who builds agentic systems and explicitly tracks model performance, described having abandoned systematic model comparison as a practical necessity, "What I originally used to do was test every single model that would come out. The iteration was slow enough that it could be done. Now I can't even keep up."

The progression of agentic AI also affected this research. As agentic capabilities became more commonplace, the separation between agentic and non-agentic AI became less obvious to interviewees. Interviewee 12 noted,

> People call things agents that you could de-

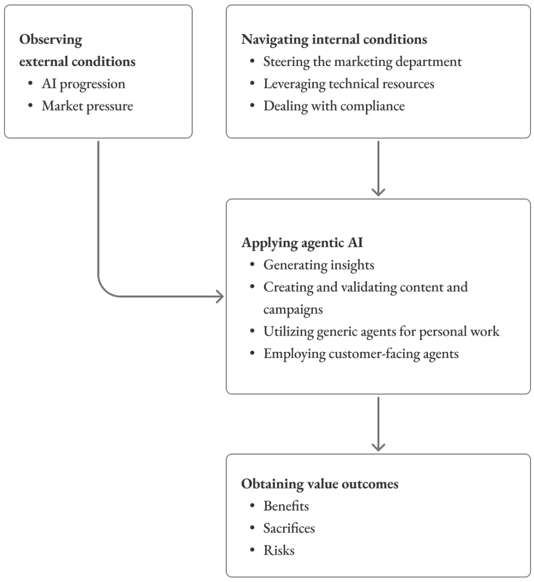

Figure 1: How managerial interventions drive value creation using agentic AI.

> bate and say, "I don't know, is that really an agent? How agentic is it?" You know, how much autonomy, how much decision-making? It, it's, it's become almost a marketing term, uh, more so than a technical definition. … You know, a lot of it comes down to it's like a continuum of being more agentic or less agentic.

While most interviewees describe the progression of AI with enthusiasm, interviewee 4 noted that it can also be overwhelming,

> It goes so fast that [employees] get overwhelmed. I think a lot of people feel that they're drowning. I think I told you, last year that I made a conscious decision not to be an expert in all the models, but in the way you implement these because you can't keep up.

The rapid progression of AI can also be seen in the rate of productization of AI. While marketers might feel like they are drowning, the release of agentic AI solutions by existing vendors can offer a safe haven. Interviewee 1 noted that "most of the thinking behind an AI strategy comes from their default supplier; it's Microsoft's AI strategy or IBM's AI strategy, because that is their source of knowledge." Interviewee 2 described a similar pattern, "We are mainly relying on the products that we use and agents that they provide us with."

Observations can also be misleading, as agentic capacity is frequently presented as operational well before it truly is. Interviewee 9 acknowledged this directly: "I think we sometimes present something a certain way while we're trying to actually make it that way in the background." Interviewee 12 described the same pattern, noting that organizations "say that you have something that you don't really have, and then you go and build it really quickly." Notably, even though this pattern was observed, none of the interviewees doubted the underlying value of agentic AI; the concern is inflated expectation. As interviewee 11 cautioned, "The big thing is hype. So AI is being overpromised all the time. And we have to reduce expectations."

4.1.2 Observing market pressure

A second external condition frequently observed by marketing managers that affects the adoption of agentic AI is market pressure. Market pressure is defined broadly as a combination of two categories: pressure from competitors and customer needs or demands. Interviewees note shifts in both categories. Regarding pressure from competitors, interviewee 9 noted how different marketing agencies are looking at each other's offerings and trying to mimic what others are doing. She noted that organizations are primarily trying "to get to parity in the present with what they're saying they can do", to ensure "speed to market" and "staying competitive". Interviewee 13 noted the value of learning from competitors,

> Talk to your peers at conferences. I think this needs to happen more on the C-level, that they exchange ideas with their peers, because that's actually where they will understand the quickest what the benefit is, and this makes it a lot better, right?

Participants described organizations feeling compelled to demonstrate AI capability, primarily to avoid the perception of falling behind rather than to generate novel value. When it comes to competitive pressure, the dominant managerial posture is defensive rather than offensive. Interviewee 8 captured the organizational reaction to new AI initiatives, "Almost everyone's been very positive about it because they want to not be left behind, you know?"

Consumers' expectations might change as market offerings evolve, but consumers are also using AI themselves. Interviewee 12 notes that "agents of consumers" might be the development that is "most disruptive to marketing" as they are a "new intermediary" and potentially create "agentic commerce." Regarding the position of agents of consumers intermediating search, interviewee 14 observed, "the world of search is now changing. And fortunately it's not going as fast as everyone thought it would change." For marketing managers, this shift translates directly into a new operational concern: the need to be visible within AI systems rather than only within traditional search engines. Interviewee 3 described this as an active focus area for her team this year, observing that organic search and paid search are "becoming less and less relevant" as AI-powered queries grow.

Looking further ahead, some interviewees described a shift that is not yet operational at scale, but is already shaping how they think and invest. Interviewee 15, a CMO at an e-commerce company, expects consumers to let their AI agents make purchasing decisions. She notes:

> I think people are going to use this much more to come to decisions and make purchases to some extent. So people will, maybe not yet, but later they'll just give their credit card details to an LLM and it's going to buy things for them on request. Or maybe even on a schedule.

This belief was the impetus for the decision to create a model context protocol server for agents to use to purchase their products. Likewise, interviewee 14 noted a future where "agents in the background are talking to each other"; he continues to note a wide variety of tasks that these agents might perform, like purchasing licenses and responding to emails.

> There's no longer doing business with a human. But there's a bot, a project, an artifact, whatever we're going to call it, that goes to the platform. The platform therefore probably also needs to look different. And it's going to search for information, select a product, make comparisons. So that for me is the further form of agentic AI.

This scenario is mostly anticipated rather than experienced. However, it already leads organizations to make anticipatory investments.

4.2 Navigating the organizational context

Navigating the organizational context describes how managers steer or move around conditions within the organization. The organizational context can be divided into three key scopes: (1) The marketing department, which the manager can steer to create alignment with AI initiatives; (2) Technical resources, which can be leveraged to implement AI initiatives. (3) The compliance framework: this bounds the potential for AI initiatives, and marketing managers do little to influence this framework.

4.2.1 Steering the marketing department

The most consistently identified barrier to adopting agentic AI is not technical but organizational. Interviewee 14 notes: "Thirty percent of all AI trajectories is just technology; sixty to seventy percent is actually us as humans or as an organization." The main obstacle seems to lie in the size of the changes that AI brings; interviewee 12 describes change as inherently difficult,

> Change is frigging hard. It's hard for us as individuals. Generally speaking, individuals don't like to change, particularly if they've been doing something a certain way, and it works well, and it works well, and they know it, and they're comfortable. [...] That challenge gets multiplied when you start talking about an organization … you've got all multiple people, you've got politics, existing incentives, existing processes.

The manager plays a key role in this process of change. As interviewee 4 put it, "In leadership, you need to set goals. You need to inspire people to believe in what you are doing and pull the people together". Steering the marketing department involves the managerial activities that help the marketing department prepare for change and guide it through the accompanying process. Managerial behaviors that comprise the steering of the marketing department include: bringing people along, stimulating experimentation, and setting direction.

The most prevalent of these steering activities is *bringing people along* which is involved with ensuring that employees feel and are ready to adopt AI. This form of steering seems to decrease resistance and helps employees adopt AI. An important part is making people feel valued during the changes of AI, as interviewee 4 noted, "people, uh, are drinking from a fire hose. And they need to have that leadership that tells them their importance, they're part of the team, and their work matters." Interviewee 13 connected this to the process of change: "Yes, we will automate these kind of things, but that still means someone needs to look over this automation. You can learn to do this".

Some employees might have a natural tendency to leverage the potential of new technology like agentic AI. For instance, interviewee 3 noted: "I am intrinsically motivated to use it, because it makes life so much easier." However, many require additional guidance. Interviewee 9 describes that when you want to bring people along, you need to get in touch with them and become concrete about how they can create value with agentic AI.

> People don't know it, so they just go back to doing their own thing until someone puts it right in front of them and says, this is what it is and this is how you use it in your day-to-day

Interviewee 4 describes a similar pattern in which he provides peer coaching to his colleagues. He describes his… "training sessions with my coworkers, showing them the agents that are being developed, or that they themselves can develop an agent. […] It becomes something that they can quickly go back to and work efficiently with."

While education and training are important components of bringing employees along, it is also vital at different levels of the organization. Interviewee 10 described a three-level model:

> Help people learn what AI can do […] leadership, middle management, and employees, each with different focus. Leadership must understand where AI is heading to create strategic horizons. Middle management must learn how to facilitate and structure experimentation. Employees must learn that AI is more than writing emails.

Another aspect in which they steer the marketing department is by *stimulating experimentation,* which includes incentivizing and supporting learning through trial and error. This behavior is primarily a response to the need to act amid uncertainty. Interviewee 12 was direct: "You have to act. You can't just be in a mode of studying this or writing memos about it." Interviewee 4 observed the cost: "We spend so much time debating which shot we have to take that that is already some cost."

Managers respond by creating structured opportunities to learn by doing. Interviewee 6 summarized the orientation: "I'd rather have something live and test and improve than be two years down the road and then finally have something live," and Interviewee 15 echoed this bias toward action: "we're more of the just-do-it type than mitigating all risks in advance." Interviewee 12 described that it can be highly beneficial to have a safe space for experimentation, "the factory where the trains are running on time" alongside "the laboratory, where you're in a position to experiment and try things and have things fail and break."

Crucially, this experimentation depends on leadership backing. Interviewee 13 described enabling it at the team level: "Provide freedom to experiment and don't set this atmosphere of fear that AI is going to replace you." Sustaining it requires endorsement from above, because, as Interviewee 12 noted, early learnings "don't always provide immediate returns" and so must be deliberately supported. Interviewee 8 described how decisive that backing was: "She just let me go and run […] the buy-in of my immediate management is pivotal to what I do." Such backing can also unlock resources directly; Interviewee 8 recounted leadership absorbing AI spend so that "anyone within the corporation can use it for free."

A third aspect of steering the marketing team is *setting direction*: naming where the team is going with AI and giving that direction force. The first part is clarity: describing where the organization is heading, which use cases matter, and what success looks like. Interviewee 6 captured the posture: "This is the vision. This is where we're going. Help me get there." Interviewee 1 noted the cost of its absence: "If you don't have clarity, jumping into execution will only surface a lot of problems." Interviewee 10 located where clarity most often breaks down: "Middle management is often forgotten, but they receive strategic targets and must operationalize them. That's where much of the friction sits."

Frequently, the impetus for the strategic direction of agentic AI is provided by an AI champion, who is typically a manager or technical employee backed by superiors. They are often driven by personal initiative. Interviewee 9 described how disproportionately impactful such a person can be: an internal platform "created originally by … our director that dabbles in development … has blown up, and so now there's millions of dollars of funding." Interviewee 8, who fits the description of an AI champion, describes his approach: "There's a new thing out, it takes me five seconds to test it, and then I know if it's ready for something that we might want to deploy. And at that point, I try to get buy-in." Champions typically combined technological know-how with business acumen, enabling them to produce examples that set direction and earn senior support.

4.2.2 Leveraging technical resources

The second scope of internal conditions concerns the technical resources on which successful implementation and use of agentic AI depends. *Leveraging technical resources* concerns managerial behaviors that use technical resources such as data, tooling, data infrastructure, and technical talent to implement agentic AI systems or improve their functionality. The use of these resources is typically described implicitly in conversations. However, when resources were not available to be leveraged, they were typically described as barriers to agentic AI initiatives. While some of these resources might reside in the marketing department, many do not; that distinction shapes how much the manager can do about them. Interviewee 6 captured the resulting relationship: the chief marketing officer (CMO) should be "the CTO's best friend," because the chief technology officer's (CTO) cooperation is necessary to execute her vision.

When resource conditions are available, managers can respond by leveraging them to pursue initiatives with agentic AI. However, critically, the data shows no clear patterns of trying to alter these conditions. The limiting factor might be that technical resources often stem from other teams outside the marketing department's direct control.

Leveraging data and technical infrastructure emerged as the primary technical prerequisites. This ranges from having data to ensuring it is accessible in the right systems, a process typically facilitated by technical infrastructure. Without accessible data and tools, agentic systems cannot act within those tools or interpret the data, limiting the ability of agentic AI to act within the organization. Interviewee 10 noted the impact of this limitation, "Technology strongly defines what can and cannot be done in marketing. Your MarTech stack bounds what personalization you can do". Interviewee 2 notes that a lack of these capabilities is the core reason his organization is behind in AI adoption, stating, "We are just in the first steps towards [AI] maturity, and the core reason is: the infrastructure is not yet in place, and the workload is really big." Interviewee 10 laments how far many organizations still are from this foundation: "How badly we still handle data. We keep saying data isn't in order and it's hard."

However, when technical resources are available, they help accelerate the process. Interviewee 11 put his organization's smooth adoption of agentic AI down to prior experience with data processing and classical AI, noting that "we have been invested in our data operations for a long time" and that "a lot of the data people already did a lot of deep learning … [which] helped us a lot."

However, interviewee 13 notes that you do not always need advanced data infrastructure or technical skills to access data with agentic AI. He notes,

> You can just dump your data …, not even dump your data. You just tell the agent, 'I have data in this source, this source, and that source. Please look at it'. Then this agent grabs all the data from different tools.

Marketing managers might also *leverage technical talent;* this can refer to letting employees with the right capabilities build, implement, or integrate agentic AI, but it can also refer to letting employees with the right abilities shape the required context for agentic AI. Here, the data shows more explicit mentions of the lack of this resource than of its utilization. Interviewee 6 was direct: "we don't have the skills yet to actually do it." Interviewee 2 described an organization with "a lot of people" overall but "not a lot of people working on the digital [marketing] landscape." Interviewee 13 argued that this gap, more than the technology itself, is where initiatives fail: "you buy gen AI, you buy a big tool, you have no one taking care of that," Interviewee 5 captured how the bottleneck has shifted: where the technology "used to always be very cumbersome … the most expensive part," the limiting factor is now "having the right people … available."

Notably, the need for talent goes beyond pure engineering. Experts describe a need for a key analytical capability: the ability to analyze the process and systems in play that might be automated with agentic AI. Interviewee 13 described this as the root cause behind many apparent AI failures:

> The stupidest example is … people want to understand their marketing data better. And it's really just like you have 20 different sources, you don't know how to connect them, you don't know how to put this all together. … The culprit is really still they don't understand what the issues are.

Similarly, interviewee 10 notes,

> In many marketing teams it's very clear what the output is, while the way people get to that output is poorly refined … the marketing manager has no idea how people get to the output. So how on earth can you then decide what to automate?"

Interviewee 1 noted why marketers might struggle to analyze these processes, noting they might find "perceived value in sitting on complicated processes" and "find comfort in bespoke engagements" with technology. She adds:

> I think there is a component that is science, which could be automated, and there's a component that is art, which is the creative and whatever. But you need to decouple the art from the science and use technology for what it's worth.

Interviewee 1 added to this that this type of deep process knowledge (a feel for system logic) is precisely what is hardest to augment with AI, because it requires a level of structured thinking that many teams have not yet developed.

When technical talent is lacking, the most common response is to source it externally. Interviewee 6 was blunt about the alternative, "using experts externally, because if we would've waited for people internally … it would've been a major blocker," the internal route meaning "I'm going to wait two years before I get space on the backlog." Interviewee 13 described the same move, in which internal teams were overwhelmed: "we are now really building a big clean room with another company because the internal IT team is completely overwhelmed." Yet the interviewees were clear that the bypass carries a cost. Interviewee 13 warned that while "agencies give you a kickstart," the real danger is that "too many companies make themselves dependent on external workforce," leaving "no one internally who understands how this is working"; the knowledge ends up held by freelancers rather than the organization. Routing around the bottleneck buys speed now at the risk of dependence and stalled capability later.

4.2.3 Dealing with compliance

A final steering behavior managers engage in is dealing with compliance. This concerns accepting or circumventing the internal compliance policies typically set by the legal and IT departments. Marketing managers seem to accept the compliance frameworks as they are; no attempts to alter or negotiate compliance policies were observed. However, they do note it as a significant barrier that they have to deal with. Interviewee 13 observed a pattern in large enterprises:

> They [employees in enterprises] are not allowed to use agentic AI a lot … there's this weird distinction people make. It's totally fine to buy any kind of tool you don't know anything about and put it [data] into the system. But agentic AI? That's completely out of the question.

This can also indirectly affect AI adoption by limiting the collection or exploitation of vital resources, such as data. Interviewee 13 describes "Companies are already afraid to use certain tools because 'how does it look on the GDPR side?', 'are we allowed to track data from our customers?' ". Interviewee 10 notes the impact of restrictive compliance policies on the culture: "First everyone does bring-your-own-AI; then someone says 'this is maybe not such a good idea' and suddenly everything is banned. Then the most enthusiastic people either disengage or leave". Interviewee 6 noted that restrictive compliance policies can prevent organizations from learning: "If we would've been stopped at the beginning, 'You cannot use this system because you're already on this system', I think that is something that's really holding people back."

Marketing teams sometimes work around the compliance policies, or ignore them altogether. Interviewee 15 noted, "If your company says it's not allowed, well then I would still do it privately … not with your company's data or insights, but think up your own things and go experiment with it" Interviewee 13 noted how shadow IT, i.e., software that is used without IT's approval, is used to start using agentic AI, where "Shadow IT can be a very good lighthouse project to make people understand what is possible."

These elements of the organizational context are consistently described as important enablers of or barriers to the effective use or implementation of agentic AI. When the right environment is in place and the AI initiatives within the marketing department are effectively steered, applying agentic AI becomes easier and more beneficial.

4.3 Applying agentic AI

Applying agentic AI refers to using or implementing agentic AI systems within the organization. Key use cases used within the marketing department are analyzed, including generating insights, creating and validating content and campaigns, utilizing generic agents for personal work, and employing customer-facing agents.

4.3.1 Generating insights

Marketers can use agentic AI to *gather insights*, which is defined as performing analytics on any form of information, including both structured and unstructured data and internal and public sources. Interviewee 3 noted, "Analytics is the place where we use [agentic AI] the most." Interviewee 15 added regarding this use case, "It does start with those insights, of course. And those insights you can then use to elevate your marketing automation to another level." Others emphasized the capability to ingest unstructured data; interviewee 11 discussed doing "analytics at scale"; he noted that it used to be the case that "a consultant has to plow through a lot of [data regarding] competitors and a lot of content, and now an AI can do most of that work".

A wide variety of scenarios in which marketers might leverage agents to generate insights were observed, including analyzing campaign performance, customer journeys, and market research. Yet users' experiences applying agentic AI to generate insights diverged depending on the available data and the integration with data sources. Interviewee 4 reported highly beneficial results after connecting an agent to live data sources, turning reporting that once occupied a "giant analytics team" into a prompt answered, "within ten minutes." Interviewee 3, pursuing the identical goal, met the opposite experience: her tool "can only digest so much data [...] if I try to dump anything over 30 [megabytes], it [...] doesn't do it." This forces manual workarounds: she notes "I pull the data from [the warehouse] myself, dump it in, and then have it analyzed." This was required because the organization's data layer was, in her words, "not there yet." The models were equivalent; the connectivity was not.

4.3.2 Creating and validating content and campaigns

Marketing teams can also use AI agents to engage in *creating and validating content and campaigns*, this refers to the entire process from ideation to publishing the full campaign, but also includes the technical elements of the campaigns like segmentation and targeting. While non-agentic AI can be used to create content, the agentic capabilities play a vital role in validating content and orchestrating the surrounding process. Interviewee 6 describes that they build a "modular system" for content creation, deployment, and evaluation:

> AI makes sure that the right elements [copy assets] are matched together and are delivered, and then it's checking which one is performing best in terms of full asset scope. And then it creates new ones [...] and then delivers that again to an audience that it thinks it's most relevant for.

Interviewee 9 describes a similar system that is focused on deploying campaigns across different cultures:

> We have something called [redacted], which is a custom journey tool, and it's basically a large language model where you constantly build up cultural insights. [...] So it can push through the end-to-end process with minimal human interaction. [...] [redacted] takes those cultural insights and creates a brief. The brief then creates a content matrix.

Regarding content creation, interviewee 11 notes the ease with which you can scale output when using AI given the right circumstances "about a thousand web pages fully automated" for one client and judged it a success, "in that use case, it worked. Because it was very standardized." Confronted with creative, brand-led work, the same operator using the same tools reached the opposite verdict: "once you have to have more creativity in the output and also if it needs to be on brand [...] AI still falls very, very short."

Interviewee 12 also noted the ability of agents to segment and target users, "There's a bunch of agents here for actually creating campaigns for marketers, doing audience segmentation for marketers, audience assembly, a lot of things." He notes that this can be done with "a composable CDP on top of cloud data warehouses like Databricks and Snowflake." Noting that this composable CDP has "a very agentic approach".

4.3.3 Utilizing generic agents for personal work

Marketing teams might also be *utilizing generic agents for personal work*, which refers to the use of general-purpose agentic systems (e.g. Claude or ChatGPT) to improve personal productivity or automate manual processes. Interviewee 4 notes how he configured an off-the-shelf agent as "personal operating system",

> [Claude] CoWork can help me build out my workflows and operationalize all the things that I do, to get more efficiency out of myself. [...] I just migrated beginning of April to Cloud Code, with an Obsidian database, with all my files on GitHub, and Notion as a human interface. [...] My productivity just went 10X.

Many use cases seem to be applicable to any type of knowledge work and might not be more or less relevant to the marketing domain. For instance, interviewee 13 describes using agents to prepare for meetings: "Claude goes into my transcription tool, gets all the different transcriptions of meetings for this one customer and starts writing down what I should talk about."

A specific use of these generic agentic systems is process automation. Often participants do not hand the entire job to an agent; they segment a process and insert AI where they trust performance. Interviewee 10 set out the model plainly: "I see process steps inside workflows. If work is a stack of workflows, then most processes can be split into six or eight steps, and in some of those steps AI is used."

Likewise, interviewee 13 described how they automate processes using generic tools with agentic capabilities: "We used Make and n8n for, like, automating all the workflows. ... You say, 'I get this data from this tool. If this happens, do this.' ... And this substituted at least for two or three people."

4.3.4 Employing customer-facing agents

Marketing teams might also be *employing customer-facing agents, which refers to letting* consumers communicate directly with agents of the organization. Applications where this occurs include customer-service and sales-assistant agents that handle routine queries and recognize commercial moments (e.g. upsells, membership changes, and re-orders).

Interviewee 12 noted that these agents have, over the past two years, crossed a quality threshold they previously never reached: "they've actually gotten really good because they're leveraging the conversational capabilities of these LLMs."

Interviewee 6 described a chatbot program that combines customer-experience improvement with direct commercial outcomes. She notes that this chatbot is built around the idea of "doing everything right for the customers... ultimately they [customers] will spend more with us, and we'll be more commercially successful". The same agent that handles a membership question can identify an upgrade opportunity and execute it. She notes that she was able to scale this project rapidly because she could expand upon an existing pre-LLM chatbot and because she had "full trust from my board," because "it was really because we had commercial success first."

Interviewee 15 took a different approach; instead of offering the agent to a customer, she now offers support for her customers' agents. She notes that this approach is still in the exploratory phase, "the value there now is actually the experimenting and staying ahead of that wave." She describes the solution,

> We have built a connector for one of our web shops that you can connect to Claude, and tell Claude: Fetch all my orders from [redacted] or I want to reorder that order. And then via that connector it actually re-orders that order.

Interviewee 12 describes this conceptually: "I think this is actually the most disruptive to marketing; this is what I would call agents of customers." He adds, "It's going to be a huge disruption because it is almost like we have got a new intermediary between us as seller and the human buyer." He wonders, "How do we engage with this generation of agents of customers?"

4.4 Obtaining value outcomes

Applying agentic AI creates outcomes for the organization. These are analyzed from a value perspective. Woodside et al. (2008) note that value has a benefits and costs component. *Benefits* concern features of the outcome that benefit the team, organization, or customer. To provide an even sharper overview of the marketing managers' experience, costs are also split into two categories: sacrifices and risks. *Sacrifices* are defined as aspects of an outcome that are negative to the organization and are highly likely to materialize; *risks* are defined as the potential aspects of an outcome that would be negative to the organization but where the likelihood of the outcome is unlikely or highly unknown.

4.4.1 Benefits

Benefits from agentic AI in marketing cluster around four themes: efficiency, scalability, extension of personal skillsets, and improvements in the quality of output. When asked about the benefits of agentic AI, interviewee 3 notes "It is so much faster, so much easier [...] essentially, speed, accuracy, depth of insights, those are really helpful." Some participants put a figure on the overall return: Interviewee 6, asked to estimate the ROI of her agentic initiatives, settled on "Let's say 10X. Because it's not two and it's also not 200," noting that the business case "is so easy to make that I'm happy to spend a bit more everywhere."

Efficiency, defined as the ability to perform a unit of work with less effort, forms the largest and most consistently described benefit. This might take the form of doing more with the same amount of effort. Interviewee 3 put a concrete figure on this, describing analysis that "would've taken, like, a week to put together" now turned around "within, like, two hours." Interviewee 12 notes "the efficiency thing is real." He cautions, however, that this benefit is not a durable advantage: "it's not a competitive advantage, certainly not in the long term [...] everybody is gonna get the benefit of efficiency through AI, so it's not really gonna be a way to distinguish yourself from your competitors." Similarly, interviewee 10 states "I think value sits primarily in efficiency, at least for now".

Efficiency can also materialize as speed. Interviewee 9 noted "speed to market" as a key benefit. Interviewee 12 connects this to the ability to experiment, stating: "Because it accelerates the time to creating something, and it can do it at typically a much smaller cost, you have the ability to run more experiments." Speed also has a knock-on effect on how quickly an organization can react to changes: for example, interviewee 4 set up "competition research agents that come back with every change that they made on their landing page, pricing, et cetera, so you can respond faster."

Scalability refers to the ability to do more units of work with the same amount of effort; typically, this requires the units of work to be of a similar nature and typically requires operations to be performed in bulk. Scale is the benefit most clearly impossible without AI. Interviewee 14 related scaling AI to having an army or interns, stating, “It actually works quite the same, and now we have whole armies of those super interns”. Scale is typically related to generating large amounts of content, many variations for personalization beyond what a marketer could do, or with automating processes that would take manual effort. The magnitudes described are striking. Interviewee 8, asked how his team produced advertising at scale, recalled that copy which “would take probably about a year” by hand was generated “within… a week” through an AI workflow. Interviewee 9 stressed that scale also means variation no person can match: “a human can only think in so many… variations,” whereas the system “can create all these different versions of a campaign.”

Extending the personal skillset refers to the ability to perform tasks with agentic AI that would typically lie outside of the skillset of the user. This is the benefit described with most enthusiasm across interviews. Participants repeatedly described AI enabling work they could not previously do, typically relating to technical capabilities. Interviewee 12 captured the broadening of individual capability: “it sort of expands the range of creative possibilities and capabilities that an individual can do.” Interviewee 10 described the downstream organizational effect: the social media manager “suddenly becomes a designer and app concept creator.” The same dynamic reshapes the seniority mix of work: interviewee 7 reframed the business case around AI-augmented juniors, noting that where a campaign once required two senior marketers, “I don’t need two of those guys, I need one, a junior, and he can do the same thing.”

The opposite idea is also described. AI is seen as being comparable to a junior employee for many tasks. This led interviewee 10 to warn, “you no longer develop juniors into seniors … a major risk: no new talent pipeline.”

Improvements in *quality of output* come in various forms; the agentic system produces better output than a human would otherwise provide. Interviewee 3 noted increased “depth of insights” in AI-generated reports. Interviewee 10 notes “I find the output quality absurdly good. If you are clear about what you want”, and interviewee 5 noted how they upgraded a large content library with AI, improving the quality compared to the human-made originals.

4.4.2 Sacrifices

Alongside benefits, participants describe a set of sacrifices; things that are knowingly given up or altered as a consequence of AI adoption. Two key categories arise from the data: financial costs and job displacement.

Financial costs are the total financial expenses incurred in implementing and operating AI systems. AI is not cheap. Infrastructure costs, API usage, vendor licensing, and the investment required to implement and maintain AI systems are consistently identified as real costs. Interviewee 7 noted: “AI is also very expensive. That’s something that we don’t take into consideration.” The scale of commitment can be substantial: interviewee 9 noted her agency had invested “millions in our agentic tool chain.” Several participants also cautioned that current prices understate the eventual cost; interviewee 16 observed that “a lot of AI is subsidized now” and expected that “these prices will go up” as the market matures.

Job replacement refers to the loss of jobs or drastic changes in roles. Job replacement is acknowledged across interviews as a potential downside of using agentic AI. However, while this might be a logical mirror to the increased efficiency, only interviewee 7 describes an occurrence of job loss, “What happened is we fired half the people, right? That’s terrible, I did not sleep for a week.” Interviewee 4 noted the general attitude, “people should be very afraid that their job is totally going to be different. And if you don’t change with the job, you find yourself out of a job.” In other examples, interviewees describe more or different work being picked up. Interviewee 13 notes regarding job displacement: “this might happen, that you get redundancies, but it is not the way I think companies are doing it right now.”

4.4.3 Risks

Risks differ from sacrifices in that they represent potential rather than certain negative outcomes. Three risk categories are prominent: hallucinations, security and privacy violations, and negative brand impact.

Hallucination risk refers to the chance that AI might provide incorrect answers that are not consistent with the provided information. While it is often noted, it is rarely described as an impactful issue. Interviewee 11 notes hallucination as a risk, but adds, “Hallucinations do not happen that often”. Interviewee 10 noted, “Hallucination is not your biggest blocker right now, and definitely not a reason to delay”. Interviewee 12 noted the need to stay aware of it:

> I do not think of AI as being infallible. […]
> If you don’t have an understanding of your

> domain, then just sort of blindly trust AI, I think most people would agree that is kind of a risky place to put yourself in.

While AI can give erroneous output, it is logical to compare that to human mistakes that would have occurred otherwise. Interviewee 3 noted that "Sometimes in analysis, people find data points to prove their point versus being objective about what the data is saying to you."

Interviewee 13 relates the performance of AI agents to a mediocre employee:

> You can always build something very quickly that gives you the same information as if you would have an analyst sitting there. It's not as good as if you would have a great analyst, but it's pretty much the same as if you have a mediocre one and still have to sense-check what comes out of it.

Characterization of agentic AI as equivalent to "a mediocre analyst that you still have to sense-check" implies a standard where some level of error is accepted as normal and managed through review, the same standard already applied to human junior analysts.

Security and privacy violations are also most consistently noted as risks. The concern centers on data that enters AI systems outside of what would be compliant, but might include any violation of law, policy, or user's trust regarding security and privacy. interviewee 9 noted that clients are "concerned that their PII, basically any of their customer data or business-sensitive data, is going to go into an AI tool," while interviewee 2 stressed the need for "maximum privacy and security in place," keeping data "secure" and "closed." Interviewee 13 described:

> Companies are already afraid to use certain tools because 'how does it look on the GDPR side?', 'Are we allowed to track data from our customers?' And this is the first step. I think the agentic level, taking like two or three tools where you're already not, uh, where you're already afraid to use them, and then giving it to an agent… that's, that's three, four levels more difficult.

Brand risk is also frequently noted and is concerned with the degradation of users' perception or the brand resulting from AI use. Content that is off-brand, tonally inappropriate, or perceived as inauthentic might negatively impact the organization. When asked for the risks of agentic AI, interviewee 6 noted: "It is mainly brand control, I would say. It is also going to say some stupid stuff, and that is a very important one." Interviewee 15 experienced the negative effects after using AI-generated messages; she observed, "The customer sees through that pretty quickly", she noted it was "just a bit too much of a shortcut with this use case"

Conversely organizations leverage agentic AI to counteract brand risk. Interviewee 3 explains "The branding team trained the assistant on our branding guidelines and different personas. And then that assistant tells you if this is brand-aligned". Likewise, interviewee 16 described governance agents that ensure generated content adheres to tone-of-voice and visual brand standards at scale, and notes that this agent is now in place to validate human content.

His case provides a good analysis of this paradox; in addition to AI, humans can also create content that is not brand-aligned. Interviewee 16 thinks AI might already outperform humans and decided to measure the difference:

> I'm logging a quality score based on these guidelines, and I do want to see this quality score go up. But I also want a human quality score next to it, so we can benchmark how we're doing compared to the traditional way of working, so we can prove it's better. Because otherwise it's a feeling.

The examples seem to indicate that brand risk is not a question about whether you are using AI, but rather how you are using AI.

Taken together, the benefits, sacrifices, and risks above describe the value agentic AI can produce. Their realization, however, is highly uneven: the same use case proved transformative in one organization and stalled in another. Table 5 illustrates this divergence for the three most frequently reported use cases.

5. Discussion

This study set out to understand how marketing managers create value with agentic AI. GenAI has moved rapidly from a technical novelty to a consequential development in marketing, yet empirical research into its implications within organizations remains scarce; the existing literature concentrates on generative AI (e.g. Elgheit, 2025; Grewal et al., 2025; Wahid et al., 2023). Scholars have responded by issuing explicit calls for research into agentic AI in the marketing field (Jain et al., 2024; Kim, 2025; Mogaji & Jain, 2024). This study answers those calls directly by providing a grounded empirical account of how marketing managers create value with agentic AI.

Table 5: Examples of comparable use cases with highly divergent outcomes

Use case	Where it created value	Where it fell short
Self-service analytics	Data sources were connected directly to the agent; Reporting requests were answered in minutes (Interviewees 15, 4)	The lack of integration or presence of data infrastructure yielded the need for manual workarounds (Interviewee 3)
Content and campaigns	Content created by Modular, validated, brand-governed pipelines (Interviewees 6, 9, 16)	Negative effect on brand due to missing quality validation and brand governance (Interviewee 15)
Customer-facing agents	A deployed chatbot driving retention and upsell (Interviewee 6)	Early or experimental deployments; withheld in a regulated sector (Interviewee 15)

5.1 Contributions to theory

The principal theoretical contribution of this study is a model of managerial value creation with agentic AI (Figure 1). The model offers an interpretive account of how value is constructed; it shows the impact of the organizational context and highlights the value of the managerial work, which includes observing external conditions, including the rapid progression of AI and the accompanying market pressure, which provide the impetus for change; *navigating internal conditions* shows how managers act in an organizational context that can enable or constrain agentic AI initiatives; *applying agentic AI* through specific use cases and usage strategies; and *obtaining* the *value outcomes* that result, which are understood as a portfolio of benefits, sacrifices, and risks.

The data shows that similar use cases can yield very different outcomes across organizations. Based on this analysis, the organizational context plays a vital role in the successful implementation and use of agentic AI, including, e.g., the availability of accessible data and the competencies necessary to implement or use AI. This work describes how marketing managers navigate these conditions. Navigating includes key managerial behaviors, such as bringing people along, that play a vital role in providing the necessary context for the successful implementation and use of agentic AI.

In line with earlier work on digital transformation by Enholm et al. (2022), the crucial role of technical conditions is observed, data and infrastructure; the organizational context, including culture, management support, organizational readiness, and more; and efficiency is still observed as one of the main benefits of AI. However, key differences were also observed. Where Enholm et al. (2022) still note compliance as a key enabler for adopting AI, the data suggests that compliance often has become overly restrictive and has become one of the main inhibitors of agentic AI.

Likewise, this work also aligns with and builds on the insights from Fernandez-Vidal et al. (2022). Fernandez-Vidal et al. (2022) note how digital leaders should become business leaders, including developing business know-how within the digital function. This study builds on this idea by mirroring it: marketing leaders should also become digital leaders, as the marketing function increasingly requires technical know-how and management skills. This conclusion implies that many of the lessons from Fernandez-Vidal et al. (2022) now apply to the marketing manager, including the need to master talent complexity, prioritize learning, and bridge the gap between technical and business capabilities.

Gebauer et al. (2020) work describes the digitalization paradox, which shows that digital transformations often do not yield the expected returns. This study highlights a different but related pattern, where the outcomes of investments in agentic AI differ significantly. While Gebauer et al. (2020) note elements related to the implementation of the specific solution as key traps, the data suggests that, when it comes to the success of agentic AI, the most significant factor is not the agentic AI initiative itself but rather the organizational context in which it is embedded.

Both the organizational resistance observed and the breadth of benefits when successful observed within the dataset seem to be much higher than what is signaled in previous works regarding digital transformation. The data suggests that this increased speed creates increased resistance and delaying behavior, where one participant noted that, "It goes so fast that [employees] get overwhelmed". This implies that the managerial behaviors that were important during earlier waves of digital transformation have become even more important in the age of agentic AI.

The first sub-question this paper sought to answer was "How do marketing managers identify and translate agentic AI opportunities?" The impetus for change comes from marketing managers who observe the changing landscape. The data shows two key external sources of inspiration that drive AI initiatives: the rapid progression of AI capability, including model releases and the roadmaps communicated by their software suppliers, and market pressure from competitors and, increasingly, from consumers who now have the ability to act through their own agents.

Observing is often passive; the progress of agentic AI, adoption by competitors, and the changing consumer behavior seem hard to miss. This seems to force marketing managers to develop a response even when learning and interpretive activities have rarely been applied.

Managers might follow the AI strategy of their main software vendors rather than doing the learning and interpretive work themselves. Furthermore, optics can be especially deceiving when it comes to agentic AI, as there is often a gap between the optics and communication regarding agentic AI solutions and the reality of these AI solutions; these solutions are often presented as successful even when they are still to be implemented altogether.

This study describes how the organizational context plays a vital role in the successful adoption of agentic AI. This aligns with earlier research on digital transformation (e.g. Ancillai et al., 2023). Among others, the organizational capacity for change, the presence of a clear direction, the availability of technical resources such as data, and the restrictiveness of the governance framework all affect whether agentic AI can be successfully implemented and leveraged. This analysis of the data empirically confirms and expands the organizational readiness criteria as proposed by Srivastava (2026), and, in line with her call to research, connects these factors to value creation. Furthermore, evidence is provided that aligns with Schmidt et al. (2026) concept of process debt. Furthermore, these concepts are found within the domain of marketing, and how managers respond to these contextual elements is analyzed.

These aspects of the organizational context seem to be able to prevent or accelerate the implementation of agentic AI, drive the resistance or enthusiasm that drives AI adoption, and bound the knowledge and access to data and systems needed to effectively use agentic AI.

The second question we sought to answer was "How do marketing managers adopt agentic AI within their organizations?" We observed that managers navigate the organizational context with a set of specific behaviors. Marketing managers engage in steering the marketing department, this includes a range of managerial behaviors, including *educating* when AI literacy is a key obstacle; *bringing people along* to decrease resistance and increase the chance of successful adoption; and *setting direction* to ensure that AI initiatives have sufficient support. We observe that marketing managers might leverage technical resources like data, tools and technical talent. The data suggests that marketing managers benefit from working tightly together with the technical leaders of the organization. When legal or IT compliance becomes an obstacle within the organizational context marketing managers typically do not seek to alter the compliance policies, but rather they deal with them; accepting them as they are.

Within the organizational context, marketing managers can seek to apply agentic AI. Four key use cases that marketing managers can apply are observed: generating insights, creating and validating content and campaigns, utilizing generic agents for personal work, and employing customer-facing agents.

A frequent application of agentic AI is the generation of insights. This aligns with the work of Hughes et al. (2025) who note that agentic AI might be used for data analysis to obtain insights and identify patterns. Interviewees note the breadth of analytical use cases; they span use cases requiring structured and unstructured data, private and public data-sources. Specifically marketers might apply agentic AI to perform competitor analysis, analyze marketing performance, and more. Interviewees also note that access to data from these AI systems is vital for these analytical use cases.

This study also shows how marketing teams use agentic AI to create and validate content and campaigns. This broad wording is used because it is observed that some interviewees are starting to automate the entire process from the idea for a campaign to the publishing of the campaign. Earlier work by Vaid et al. (2025) already described that GenAI can be used to segment and target customers. Likewise, the ability of GenAI to create content was also already described (e.g. Grewal et al., 2025; Wahid et al., 2023). This study adds to existing research by showing that agentic AI can create a more integrated process where the generated content can be tuned to data regarding the customer or segment of customers and can now also be validated by agentic AI.

The third use case is likely not marketing-specific, but frequently noted by the interviewees: the utilization of generic agents for personal work. These findings seem to extend the findings from Hasselwander et al. (2026) regarding the AI super assistant: Data shows that marketers are already configuring agentic AI systems to perform a wide variety of tasks. It seems like these systems are increasingly implemented without a specific task in mind but are rather employed as a general purpose assistant. Additionally, previous work also showed that agentic AI can be used to automate processes (Hughes et al., 2025). This study also shows that marketers can be in control of these automations themselves, they use tools that do not require programming experience to automate their own processes.

The last key use case observed relates to the deployment of agentic agents that can be leveraged by the consumer. The concept of consumer facing chatbots has now been around for a while (Dutta et al., 2021). However, interviewees now describe agentic systems that have the ability to act, rather than simply providing answers to questions. These agentic chatbots are typically offered by the companies and expected to act on behalf of the customer within the bounds set by the company. Some of these new agentic chatbots seem to be blurring the boundary between a customer support agent, a sales representative, and an agent that can be used as a tool to perform tasks. We also see a new trend that aligns with the concept of the AI super assistant posed by Hasselwander et al. (2026), while still in an experimental phase, companies start to expose specific capabilities through servers that can be integrated with the consumer's agentic AI system.

The model's final component concerns the value outcome that the agentic AI initiatives create. We follow the perspective of Woodside et al. (2008) and analyze value from a benefits and costs perspective. More specifically, the data shows many costs with a low likelihood of materializing, therefore we separated costs into sacrifices, i.e. costs that are expected to materialize and risks which this study defines as potential aspects of an outcome that would be negative to the organization but where the likelihood of the outcome is unlikely or highly unknown.

Almost all of the specific benefits, sacrifices, and risks that are described in this study have been found in previous work. For example the benefits; efficiency, scale, and quality of output are described by Hughes et al. (2025), Kim (2025), or Brynjolfsson et al. (2025). Likewise, the sacrifices, financial costs and job displacement, and risks, hallucinations, security and privacy violations, and negative brand impact are all described by Hughes et al. (2025). This study confirms these effects of agentic AI empirically and shows that they apply to the use cases deployed within the marketing domain.

One of the benefits that we did not find in earlier works relates to the extension of personal skillsets, this perceived benefit relates to marketers being able to perform tasks that they typically are not able to do. Interviewees see this pattern and connect it to the cost of job displacement. However, different interviewees interpret this differently, some note that juniors might become less relevant, noting that with agentic AI you essentially already have an army of juniors. Others note that it is instead the seniors that might get displaced because a junior with AI can do the same work while costing less.

This study also adds perspective to the explored benefits, sacrifices and risks by analyzing how they are mediated by managerial behaviors, and how specific risks can be mitigated by applying agentic AI effectively. Table 5 shows examples of specific use-cases and describes when they provided value and why they did not provide similar value in other scenarios. The organizational context and managerial behaviors in general stand out in their role for value creation.

But additionally a few specific patterns concerning risks warrant analysis. First, hallucination is often mentioned as a key risk; however, some interviewees mention the risk of human failure as well. Some noted that agentic AI might be less likely to make mistakes, but none indicated to have measured and compared the failure rates of agentic AI and humans. Similarly, an often mentioned risk of agentic AI is brand risk; here we observe a similar pattern, while many mention content that does not meet the quality bar or does not align with the brand as a risk, some note that they use AI to validate quality and alignment with the brand. These agentic AI content validation systems might be deployed to validate content generated by both human employees and AI systems and in doing so agentic AI might significantly lower brand risk compared to a situation where AI is not deployed.

5.2 Practical implications

This study describes how marketing managers create value using agentic AI, and a key element in doing so is navigating the organizational context. The data shows some practical implications that can provide guidance to managers who seek to benefit from agentic AI. In addition to the points made below, Appendix B contains a list of practical suggestions that stem from this research and are directly applicable for a marketing manager.

Our study indicates that the organizational context plays a vital role in the process of creating value using agentic AI. This implies that investments in setting up the right internal conditions to help ease the implementation of agentic AI seem sensible. Some elements have already been mentioned explicitly, others can be derived from the implicit patterns in the data.

The condition of the marketing team is important; interviewees note the value of bringing people along; employees can feel overwhelmed by the rapid changes that AI brings, but if you show them that you value them even though things are changing and if you help them master the skills necessary to work in the new context you can ease the transition. As interviewee 13 noted: "Yes, we will automate these kind of things, but that still means someone needs to look over this automation. You can learn to do this". This is connected to the need for education, which is described by many as critical to help employees apply agentic AI effectively.

An interesting interplay arises between compliance and experimentation. Since the possibilities of agentic AI are progressing so fast, it can be hard to know what works and what does not. Experimentation helps organizations to act fast, and learn while doing. However, we also note that compliance policies can be a big obstacle to implementing agentic AI within the organization. Interviewee 12 notes a possible solution: "the laboratory", a safe space where employees can experiment with AI. The data suggests that marketing managers deal with compliance policies as they are. We advise marketing managers to negotiate with legal and IT to set up these laboratories in a compliant manner, ensuring space for experimentation without breaking compliance policies.

The study observes a further blurring between the needs and responsibilities of technical leaders and marketing leaders. Where Fernandez-Vidal et al. (2022) note that digital leaders should become business leaders, we note that marketing leaders should also become digital leaders, as the marketing function increasingly requires technical know-how and management skills. Interviewee 6 noted that the CMO should be "The CTO's best friend," noting that the CTO's help is necessary to execute her vision, and noted that she used external technical capacity because the speed of the internal technical team was not sufficient for her roadmap. It might be advisable for marketing managers to take a more active role in the development and allocation of technical resources, to ensure that the necessary data, tooling, data-infrastructure and technical talent is available to support the needs of the marketing team, especially due to their centrality to agentic AI initiatives. The marketing manager might be able to help set direction, negotiate resources, or influence infrastructure decisions to align with the needs of the marketing department.

5.3 Limitations & future research

A limitation of this study lies in the speed at which AI is evolving. Many insights regarding the role of the manager align with earlier works on transformations (e.g. Fernandez-Vidal et al., 2022) and are likely to be relevant toward the future; others like the impact of specific use cases might be heavily dependent on the temporal situation. Furthermore, as this research applies a qualitative approach, the outcomes of this study could be transferable to other contexts (Shenton, 2004), most notably our dataset is skewed towards the Netherlands, so the results might not be transferable to other markets, for instance due to a different cultural or legislative climate. To balance these limitations, future research could make comparisons across industries and conduct longitudinal research to monitor the evolving response of organizations and consumers to the (also) evolving nature of AI on value creation with agentic AI and the role of the marketing manager therein.

Another limitation lies in the gap between optics and reality; while many interviewees were open about this gap, participants might perceive a benefit in describing their initiatives as successful. This study relies exclusively on interview data and is based on an inductive approach. In line with Polsa (2013), research could use different data collection methods and leverage different research approaches. For example, qualitatively measuring the effects of specific agentic AI solutions on specific values, like time saved, should provide an important balancing force to the grounded theory approach used in this study.

Finally, due to both the emerging nature of agentic AI and the exploratory nature of this study, we call for research to study specific elements in more depth. For example, this research does not go into detail regarding how specific technical features or configurations of agentic AI systems might affect business outcomes; this research might provide another lens on the high variance regarding value outcomes from similar use cases. Additionally, research could make a more detailed analysis of the specific behaviors of marketing managers, analyzing the components of behaviors like "bringing people along" and their impact on the outcomes within the organization.

6. Conclusion

This study asked how marketing managers create value with agentic AI. Drawing on seventeen interviews with marketing managers and AI experts, analyzed through constructivist grounded theory, the findings propose a model that describes how the marketing manager creates value using agentic AI (see Figure 1) containing the managerial behaviors: observing external conditions, navigating the organizational context, applying agentic AI, which together result in obtaining value outcomes.

Empirically, this study answers recent calls for research into agentic AI in marketing (Jain et al., 2024; Kim, 2025; Mogaji & Jain, 2024): it provides one of the first grounded accounts of how marketing managers create value with agentic AI in practice.

Our result shows the importance of the organizational context to the value created by agentic AI and describes how marketing managers navigate this context. Additionally, we describe the value of agentic AI for marketing and note a list of common use cases. Furthermore, we provide practical advice that marketing managers can apply within their organization.

While the use of agentic AI is likely to continue evolving rapidly, the higher-order principles that describe managerial responses to this ever-improving technological field are likely to be more durable. For marketing managers living, as one participant put it, through both "an era of change" and "a change of era," the practical message coincides with the theoretical one: the value of agentic AI will be determined not only by the agentic AI initiative itself but mainly by the organizational context in which this initiative is embedded.

7. References

Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey. *IEEE Access*, *13*, 18912–18936.

Almquist, E., Senior, J., & Bloch, N. (2016). The elements of value. *Harvard Business Review*, *94*(9), 47–53.

Ancillai, C., Sabatini, A., Gatti, M., & Perna, A. (2023). Digital technology and business model innovation: A systematic literature review and future research agenda. *Technological Forecasting and Social Change*, *188*, 122307.

Anthropic. (2024). *Introducing the model context protocol.*

Anthropic. (2025). *Equipping agents for the real world with agent skills.*

Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The rise of agentic ai: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, *17*(9), 404.

Bauer, J., & Jannach, D. (2021). Improved customer lifetime value prediction with sequence-to-sequence learning and feature-based models. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, *15*(5), 1–37.

Blumer, H. (1986). *Symbolic interactionism: Perspective and method.* Univ of California Press.

Brinker, S., & Riemersma, F. (2025). *The state of martech 2025.* Chiefmartec, MartechTribe.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, *33*, 1877–1901.

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, *140*(2), 889–942.

Burkhardt, S., & Rieder, B. (2024). Foundation models are platform models: Prompting and the political economy of ai. *Big Data & Society*, *11*(2), 20539517241247839.

Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis.* Sage.

Charmaz, K. (2014). *Constructing grounded theory (2nd ed.)* Sage.

ChatGPT. (2026). *Skills in chatgpt.*

Claude. (2025). *Introducing agent skills.*

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T., & Owen, D. (2024). The rising costs of training frontier ai models. *arXiv preprint arXiv:2405.21015.*

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding.

Duarte, V., Zuniga-Jara, S., & Contreras, S. (2022). Machine learning and marketing: A systematic literature review. *IEEE Access*, *10*, 93273–93288.

Dutta, S., Assem, H., & Burgin, E. (2021). Sequence-to-sequence learning on keywords for efficient faq retrieval. *arXiv preprint arXiv:2108.10019.*

Elgheit, E. A. (2025). Generative ai as a disruptive innovation: Implications for marketing strategic transformations.

Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, *24*(5), 1709–1734.

Fernandez-Vidal, J., Perotti, F. A., Gonzalez, R., & Gasco, J. (2022). Managing digital transformation: The view from the top. *Journal of Business Research*, *152*, 29–41.

Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997.*

Gebauer, H., Fleisch, E., Lamprecht, C., & Wortmann, F. (2020). Growth paths for overcoming the digitalization paradox. *Business Horizons*, *63*(3), 313–323.

Grewal, D., Satornino, C. B., Davenport, T., & Guha, A. (2025). How generative ai is shaping the future of marketing. *Journal of the Academy of Marketing Science*, *53*(3), 702–722.

Guercini, S. (2022). Scope of heuristics and digitalization: The case of marketing automation. *Mind & Society*, *21*(2), 151–164.

Guercini, S. (2023). Marketing automation and the scope of marketers heuristics. *Management Decision*, *61*(13), 295–320.

Gupta, S., Lehmann, D. R., & Stuart, J. A. (2004). Valuing customers. *Journal of Marketing Research*, *41*(1), 7–18.

Hasselwander, M., Sunio, V., Lah, O., & Mogaji, E. (2026). Toward agentic ai: User acceptance of a deeply personalized ai super assistant (aisa). *Journal of Retailing and Consumer Services*, *89*, 104620.

Heimbach, I., Kostyra, D. S., & Hinz, O. (2015). Marketing automation. *Business & Information Systems Engineering*, *57*(2), 129–133.

Holmström, J. (2022). From ai to digital transformation: The ai readiness framework. *Business horizons*, *65*(3), 329–339.

Hughes, L., Dwivedi, Y. K., Malik, T., Shawosh, M., Albashrawi, M. A., Jeon, I., & Walton, P. (2025). Ai agents and agentic systems: A multi-expert analysis. *Journal of Computer Information Systems*, *65*(4), 489–517.

Jain, V., Wadhwani, K., & Eastman, J. K. (2024). Artificial intelligence consumer behavior: A hybrid review and research agenda. *Journal of Consumer Behaviour*, *23*(2), 676–697.

Järvinen, J., & Taiminen, H. (2016). Harnessing marketing automation for b2b content marketing. *Industrial Marketing Management*, *54*, 164–175.

Kastner, J. K., & Hong, S. J. (1984). A review of expert systems. *European Journal of Operational Research*, *18*(3), 285–292.

Kim, J. (2025). Advertising in the age of agentic ai: Call for research. *Journal of Interactive Advertising*, *25*(3), 215–221.

Kohavi, R., Longbotham, R., Sommerfield, D., & Henne, R. M. (2009). Controlled experiments on the web: Survey and practical guide. *Data Mining and Knowledge Discovery*, *18*(1), 140–181.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, *33*, 9459–9474.

Li, M., Guo, S., & Liu, R. (2025). Bert-based consumer sentiment analysis for personalized marketing strategies.

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage.

Little, J. D. (1979). Decision support systems for marketing managers. *Journal of Marketing*, *43*(3), 9–26.

Mayer, H., Yee, L., Chui, M., & Roberts, R. (2025). *Ai in the workplace: A report for 2025*. McKinsey & Company.

Merriam, S. B. (1998). *Qualitative research and case study applications in education*. Revised; Expanded from "Case Study Research in Education" .Jossey-Bass Publishers, 350 Sansome St, San Francisco, CA 94104.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality.

Mitchell, A. A., Russo, J. E., & Wittink, D. R. (1991). Issues in the development and use of expert systems for marketing decisions. *International Journal of Research in Marketing*, *8*(1), 41–50.

Mogaji, E., & Jain, V. (2024). How generative ai is (will) change consumer behaviour: Postulating the potential impact and implications for research, practice, and policy. *Journal of Consumer Behaviour*, *23*(5), 2379–2389.

Moralles, C., Costa, L. A. L., Rigo, S. J., Kunst, R., Souza, V. C., Silva, E. P., & Roehrs, A. (2026). A systematic literature review of agentic ai: Definitions, architectures, and challenges.

OpenAI. (2022). *Introducing chatgpt*.

OpenAI. (2023). *Function calling and other api updates*.

Parker, C., Scott, S., & Geddes, A. (2019). *Snowball sampling*. SAGE research methods foundations.

Polsa, P. (2013). The crossover-dialog approach: The importance of multiple methods for international business. *Journal of Business Research*, *66*(3), 288–297.

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). *Improving language understanding by generative pre-training*. OpenAI.

Schafer, J. B., Konstan, J. A., & Riedl, J. (2001). E-commerce recommendation applications. *Data Mining and Knowledge Discovery*, *5*(1), 115–153.

Schmidt, R., Alt, R., & Zimmermann, A. (2026). Agentic ai readiness: A process-oriented assessment framework.

Shenton, A. K. (2004). Strategies for ensuring trustworthiness in qualitative research projects. *Education for Information*, *22*(2), 63–75.

Shinde, P. P., & Shah, S. (2018). A review of machine learning and deep learning applications.

Srivastava, R. (2026). *From ai hype to agentic reality: A readiness lens for sustainable enterprise adoption*.

Thapa, P., Khadka, S., Sharma, P., Rai, A., Bhattarai, U., Regmi, K., & Bohara, P. (2026). Agentic ai for autonomous marketing campaign execution: A comprehensive survey and open-source implementation.

Vaid, S., Puntoni, S., Honig, B., & Michael, K. (2025). In this special issue: When attention is all marketers need—artificial intelligence in marketing. *IEEE Transactions on Technology and Society*, *6*(3), 242–249.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., & Polosukhin, I. (2017). Attention is all you need.

Wahid, R., Mero, J., & Ritala, P. (2023). Written by chatgpt, illustrated by midjourney: Generative ai for content marketing. *Asia Pacific Journal of Marketing and Logistics*, *35*(8), 1813–1822.

Woodside, A. G., Golfetto, F., & Gibbert, M. (2008). *Customer value: Theory, research, and practice*. In Creating; managing superior customer value (Advances in Business Marketing; Purchasing, Vol. 14, pp. 3–25). Emerald Group Publishing.

Xu, Y., Liu, X., Cao, X., Huang, C., Liu, E., Qian, S., & Zhang, J. (2021). Artificial intelligence: A powerful paradigm for scientific research.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). React: Synergizing reasoning and acting in language models.

Appendix A. Interview Guide and Protocol

To ensure relevant topics are discussed, a list of questions and an interview protocol are utilized.

A.1. Interview Guide

A.1.1. Primary questions

1. Background of interviewee:
   - What is your role within the organization? *(key responsibilities)*
2. Openers to establish common ground:
   - What does generative AI mean to you?
   - What does agentic AI mean to you?
1. How are you already using agentic AI, if at all?
   - What are the benefits of this usage, if any?
   - What are the drawbacks of this usage, if any? *(financial, hours, negative side-effects etc.)*
   - What are the risks of this usage, if any?
   - What steps were taken to deploy Agentic AI in the organization?
   - What obstacles, if any, were faced whilst deploying Agentic AI?
   - What did you need to deploy Agentic AI? *(infrastructure, advice, other resources)*
2. What agentic AI solutions is your organization looking to implement in the near future, if any?
   - What are the benefits that you are looking for?
   - What are the drawbacks of this usage, if any?
   - What are the risks of this usage, if any?
   - What obstacles, if any, do you expect to face in implementing Agentic AI?
   - What tailwinds, if any, do you expect to face whilst implementing agentic AI?
3. (Optional) What are your strategic goals for the next year that will leverage Agentic AI?
   - What benefits do you expect to create using these initiatives?
   - What drawbacks do you expect to face during these initiatives?

A.1.2. Potential backup questions:

4. What is the impact of agentic AI on your team?
5. You mentioned *[subject x]* can you elaborate on that?
6. Did your organization experience *[element y]*?

A.2. Interview protocol

Interviews should take about 1 hour. The interviews will not strictly adhere to the questions in the guide; while the guideline is relevant to ensure that key insights will be gleaned from the interviews, room is left to go down relevant tangents to dive into what is relevant to the interviewee.

Information between brackets is not to be considered a direct part of the question but rather a reminder to the interviewer to ensure that the question is interpreted in the expected manner. The list of backup questions includes relevant questions to steer the interviewee to consider the intended scope without being leading questions. Additionally, Questions 2 and 3 in section A.1, and their follow-up questions, might be applied multiple times, since organizations can support multiple AI initiatives.

Appendix B. Putting the model to work: a manager's checklist

This appendix translates the practical implications of Section 5.2 into an operational checklist, organized by the four stages of the model developed in Chapter 4 (Figure 1): observing the external context, auditing the organizational conditions, steering the organization, applying agentic AI, and mediating the flow through to value outcomes. It is intended as a working tool. A marketing manager can use it to see which conditions are ready to support agentic AI, and what the manager can do to improve the chance that the initiative will deliver value. The items distil patterns reported across the interviews in Chapter 4.

Observing: sensing the changing context

- Watch the emerging use cases, and tools that are gaining traction.
- Monitor competitors' adoption of agentic AI.
- Monitor shifts in consumer behavior, including customers acting through their own agents ("agents of customers").
- Stay wary: what you hear or read might not reflect reality.
- Learn from others: leverage events, case studies, and stories.

Auditing and steering: proactive navigation of the internal conditions

- Assess AI literacy at every level and train employees to be ready for the future.
- Audit data availability, infrastructure, and the technical talent to use them.
- Leverage agencies where internal teams fall short (including lacking the necessary speed).
- Confirm strategic direction and secure senior-leadership backing before scaling.
- Map governance, legal, and compliance constraints (such as GDPR) and see if you can create a laboratory where you can experiment without violating compliance frameworks.
- Catalogue data to ensure that it can be interpreted by AI.
- Steer what you control: educate (differentiated by level), run experiments, bring people along, provide clarity, and champion the work.
- Ensure people feel valued and know what their place is after AI initiatives are implemented.
- Do not fall into analysis paralysis; try to find a logical starting point instead. You can always pivot later.
- Identify, empower, and protect AI champions to ensure they can bring along your team.
- Fund and protect experimentation: centralize or absorb its cost so teams are not penalized for early, low-return learning.

Applying: putting agentic AI to work

- Try to define clear problems and workflows where you can leverage agentic AI
- When you do not trust the risks, monitor and keep humans in the loop.
- Match ambition to maturity across use cases (analytics, content creation and validation, generic agents, customer-facing).
- For customer-facing agents, build on an existing baseline; put brand and tone-of-voice governance in place first. Only scale after you have gotten signals that it generates net-positive value.

Mediating Value Outcomes: turning application into value

- State the value you are pursuing (efficiency, scale, skill extension, quality) and try to measure impact. Do this not just to show the value of this use case, but also to be able to unlock future budgets.
- Monitor risks and deploy agentic AI to validate quality or mitigate risk, including brand risk, security and privacy violations, and hallucinations.
- Calibrate oversight and mitigation strategies to the cost of error and the difficulty of evaluating the output; do not overdo it for small use cases.
- Decide consciously whether to follow (parity) or to differentiate (novel use cases); efficiency gains are not a durable advantage once competitors share them.
- Treat AI budgeting as a strategic act: returns can be large but uncertain, and conventional business cases tend to under-fund experimentation.