

Association for Information Systems

AIS Electronic Library (AISeL)

ECIS 2025 Proceedings

European Conference on Information Systems
(ECIS)

June 2025

How Organizations Design Portfolio Management to Govern AI: A Taxonomy Approach

Simon Sturm

University of St. Gallen, simongabriel.sturm@unisg.ch

Benjamin van Giffen

University of Liechtenstein, benjamin.vangiffen@uni.li

Follow this and additional works at: <https://aisel.aisnet.org/ecis2025>

Recommended Citation

Sturm, Simon and van Giffen, Benjamin, "How Organizations Design Portfolio Management to Govern AI: A Taxonomy Approach" (2025). *ECIS 2025 Proceedings*. 9.

https://aisel.aisnet.org/ecis2025/ai_org/ai_org/9

This material is brought to you by the European Conference on Information Systems (ECIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2025 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

HOW ORGANIZATIONS DESIGN PORTFOLIO MANAGEMENT TO GOVERN AI: A TAXONOMY APPROACH

Completed Research Paper

Simon Sturm, University of St.Gallen, Switzerland, simongabriel.sturm@unisg.ch

Benjamin van Giffen, University of Liechtenstein, Liechtenstein, benjamin.vangiffen@uni.li

Abstract

Advances in artificial intelligence (AI) create game-changing opportunities and new risks for organizations. Increasing competition and emerging regulation pressure them to govern AI for strategic use. Portfolio management has been identified as a key mechanism. However, AI is different from technologies for which portfolio management has proven effective. How organizations design portfolio management to govern AI is not well understood. Therefore, we conceptualized key design decisions in AI portfolio management (AIPM) by developing a multi-layer taxonomy. We developed our taxonomy based on five exploratory case studies. We evaluated it through expert interviews and a practitioner intervention. Our taxonomy shows that many IT portfolio management characteristics also apply to AIPM. However, AI introduces complexities that require novel design decisions. Our taxonomy captures design knowledge that enables future causal / predictive theorizing on AIPM. Practitioners can use our taxonomy as a tool to establish and improve AIPM practices.

Keywords: Artificial Intelligence, Governance, Portfolio Management, Taxonomy, Design Science

1 Introduction

Artificial intelligence (AI) technologies offer game-changing opportunities for organizations to create value. Generative AI tools such as ChatGPT and Stable Diffusion are expected to unlock an annual economic potential of \$4.4 trillion (McKinsey & Company, 2023). At the same time, AI poses new threats. Sixty-eight percent of CEOs believe that AI will increase competition in their industries (PWC, 2024). Strategic governance of AI is thus crucial (Borges et al., 2021). The main challenges lie not in the technology itself, but in managing AI for business value (Borges et al., 2021). They include deciding which AI initiatives to pursue, aligning them with business strategy, and allocating resources efficiently. Portfolio management has been proposed to address these challenges and enable the strategic use of AI (Davenport, 2018; Fontaine et al., 2019; Haefner et al., 2023).

Portfolio management is an investment concept that originated in economics (Markowitz, 1952). Today, portfolio management is common practice in information technology (IT) governance and featured in frameworks like COBIT (ISACA, 2019). Applying portfolio principles means managing IT like a financial investor would: Selecting assets that fit the strategy, combining them to balance risk and return, and optimizing the resources invested (Jeffery & Leliveld, 2004). Portfolio management has proven to be an effective approach for organizations to manage risk, improve strategic alignment and increase the value of their IT investments (Chiang & Nunez, 2013; Reyck et al., 2005). Therefore, using portfolio management to govern AI seems promising.

However, AI is different from the deterministic digital technologies for which portfolio management has proven effective. Unique AI features (autonomy, learning, inscrutability; Berente et al., 2021) make its strategic use much more complex (Borges et al., 2021). A key aspect is that data dependency makes

many AI initiatives experimental and their outcomes difficult to predict (Martínez-Plumed et al., 2021). AI thus introduces a type of uncertainty that IT portfolio management does not typically face (i.e., whether something can be done at all). This naturally raises the question of the extent to which portfolio management can accommodate the unique features of AI. However, the existing literature tends to go no further than advocating AI portfolio management (AIPM) (e.g., Fountaine et al., 2019; Haefner et al., 2023) and lacks empirical investigation. This is problematic. Because without understanding how organizations design portfolio management for AI, its potential will remain underutilized. Scholars have thus called for empirical research that details the characteristics of AI governance mechanisms like portfolio management (Birkstedt et al., 2023). We respond to these calls and ask: *Which aspects exist for organizations to design portfolio management to govern AI?*

To address this research question, we conceptualized key design decisions in AIPM by developing a multi-layer taxonomy (Kundisch et al., 2022). We developed our taxonomy based on five exploratory case studies. We found that organizations design AIPM in twelve dimensions that span three layers: *AI Portfolio Strategy*, *AI Portfolio Management Process*, and *AI Portfolio Performance Management*. We developed our taxonomy over four iterations. We refined it in three evaluation cycles, including expert interviews and a practitioner intervention. Our work contributes to the discourse on governance and strategic use of AI in organizations. Our taxonomy captures design knowledge that enables future causal/predictive theorizing on AIPM (Gregor, 2006). Practitioners can use our taxonomy as a resource to establish and improve AIPM in their organization.

2 Conceptual Background

2.1 Strategic Use of AI Technology and the Portfolio Management Concept

Advances in AI technology enable organizations to delegate certain cognitive tasks to “a distinct form of autonomous and self-learning agency” (Taddeo & Floridi, 2018, p. 751). On one hand, this capability provides great opportunities to create business value. On the other hand, it increases competitive pressures. Emerging AI regulation also poses substantial business risks. Violating the European AI Act (e.g., by using AI to predict individual behavioral patterns) could result in fines of up to €30 million or six per cent of annual global turnover. Both the potential business value and the risks are driving organizations to use AI strategically. However, most of them still struggle to move beyond exploratory stages in their AI journey, where they “let a thousand flowers bloom” (van Giffen & Ludwig, 2023b). To enable the strategic use of AI, organizations need to establish robust governance mechanisms that link AI activities with strategy (van Giffen & Ludwig, 2023a). Portfolio management has been identified as a potential solution.

The portfolio management concept originated in economic theory. Specifically, in the seminal work on modern portfolio theory (MPT; Markowitz, 1952). MPT explains how investors (should) allocate financial resources to maximize returns. MPT shows that investors can optimize the inherent risk-return trade-off by combining assets with uncorrelated risk/return characteristics (Fabozzi et al., 2002). The underlying mechanism is called diversification (Markowitz, 1952). Diversification allows investors to minimize risk without sacrificing return, or maximize return without increasing risk (Fabozzi et al., 2002). Diversification is not specific to financial assets such as stocks and bonds, and has thus been adapted to manage a variety of entities. Today, portfolio management is a key IT governance concept and part of frameworks like COBIT (ISACA, 2019). Portfolio management helps organizations to manage risk, improve strategic alignment, and maximize the value of IT investments (Reyck et al., 2005). Therefore, the idea of using portfolio management to govern AI is promising.

2.2 Need for Taxonomizing Design Decisions in AI Portfolio Management

In recent years there has been a growing interest in AI and portfolio management. However, research has tended to focus on AI *in* portfolio management rather than portfolio management *of* AI. For example, Cheng et al. (2022) investigate predictive models to optimize product-market combinations in product portfolio management. Few studies have been published on portfolio management of AI. The

literature has tended to focus on operational aspects and adopt a technical perspective. This has led authors like Bosch et al. (2021) to explore the challenges in deploying machine learning models at scale. As an AI governance mechanism, however, portfolio management is not well understood. While authors such as Haefner et al. (2023), Fountaine et al. (2019), and Davenport (2018) have identified portfolio management as a key enabler for the strategic use of AI, little is known about how organizations implement it. Research to understand how organizations design portfolio management to govern AI is needed for two reasons. First, it is unclear how established IT governance concepts like portfolio management need to be adapted to accommodate AI's unique features (Birkstedt et al., 2023). Second, there is probably neither a single nor "best" way to approach AIPM. Understanding the options could thus help organizations to design AIPM configurations that meet their needs. The goal of our research is to address this issue. To this end, theories that describe/analyze how organizations design portfolio management to govern AI are particularly valuable (Gregor, 2006). Taxonomy development is a common way to engage in this type of analytic theorizing and create the basis to generate causal/predictive knowledge in the future (Gregor, 2006).

Taxonomies are structuring artifacts (Nickerson et al., 2013). They are built to differentiate objects based on characteristic similarities and differences (Nickerson et al. 2013). Taxonomies enable researchers and practitioners to describe and analyze the core components of a phenomenon (Nickerson et al. 2013). Since we know much more about the strategic potential of AIPM than about how organizations (choose to) implement it, clarifying key design decisions is crucial.

3 Taxonomy Design Research Method

We developed our AIPM taxonomy following Kundisch et al. (2022). We chose this method because it extends the well-established method for taxonomy building by Nickerson et al. (2013) and also offers guidance on their evaluation. We will detail the build and evaluation activities in the following.

3.1 Validating the Phenomenon and Specifying the Design Project

We engaged with practitioners before developing our taxonomy. We asked for their views on how AI differs from IT portfolio management. We also asked if they saw value in conceptualizing design decisions in AIPM across organizations and industries. We initiated the taxonomy development when practitioners validated both aspects. As an AI portfolio manager explained: *"We have already talked about the characteristics of AI and IT projects. I think it follows that it is good to treat them separately. [...] An AI project is much closer to innovation. You have more dependencies, especially in terms of data, and a much more iterative process. [...] You can also fail. Probably more often than in a traditional IT project. [...] So it's a very risky investment. [...] There is no established standard. Everyone builds it [AIPM] themselves based on their own experiences and problems. [...] There is no book I could buy with best practices. I haven't even found anything useful about portfolio management in general."* Similarly, a Chief AI Officer stated: *"In traditional IT portfolio management, some aspects that are highly relevant to AI don't really play a major role. [...] I think the sheer speed of AI innovation alone makes it reasonable to treat it separately. [...] Our IT portfolio process also makes it difficult to really drive change. It tends to include only initiatives that are already widely accepted and supported. But in the context of AI innovation, [...] I can't wait for the whole organization to give me the green light"*.

We then specified our taxonomy design project in three dimensions (Kundisch et al., 2022). **(1) Users and Anticipated Use:** We design our taxonomy for academics and practitioners in the IS field. We anticipate that researchers may use it to advance the current understanding of AI governance in organizations. The taxonomy will provide a classification scheme that may be useful to describe and analyze how organizations design portfolio management to accommodate unique AI features. The taxonomy may thus help to situate, or guide future empirical research on the emerging phenomenon. We envision that practitioners, particularly those in strategic AI roles, may use the taxonomy for decision-making. The dimensions and characteristics outline design options that they may find useful to create a shared understanding of, establish, or improve AIPM practices. **(2) Meta-Characteristic:** We chose design decisions as the meta characteristic for our AIPM taxonomy. That means we assume that

deliberate choices drive observable differences in organizational AIPM design. We thus require evidence that a characteristic does neither occur necessarily (lacking observable variability) nor randomly (lacking an indicator of deliberate choice) to consider it in our taxonomy. **(3) Ending Conditions & Evaluation Goals:** We defined six objective and four subjective conditions to conclude the taxonomy building (Nickerson et al., 2013). We chose completeness, robustness, simplicity, and understandability as subjective ending conditions, as these taxonomy properties are most frequently tested in the information systems field (Kundisch et al., 2022).

3.2 Collecting Data Through Exploratory Case Studies

Since AIPM is an emerging phenomenon, we conducted five exploratory case studies to collect empirical data (Käss et al., 2024). We chose this approach because our initial engagement with practitioners indicated that AIPM is both a complex and time-sensitive issue, so rapid discussion of research findings would be particularly beneficial (Käss et al., 2024). We purposefully selected the cases to generate the greatest information richness feasible within our study. We based our sampling criteria on key factors shaping how organizations design IT governance systems, as listed in the COBIT framework (ISACA, 2019). Specifically, we sought variation in industry, location, compliance requirements, AI sourcing model, and AIPM implementation (Figure 1). We included organizations that implemented AIPM and others that intended or were interested in doing so. This is because understanding initial considerations and expectations can also yield valuable design knowledge.

Case	Industry	Number of Employees	Location	Compliance Requirements	AI Sourcing Model	AI Portfolio Management	Data Collection		
							Document Analysis	Expert Interviews	Participant Observation
Alpha (C1)	Insurance	> 100.000	Germany	Normal	Hybrid	Implemented	●	●	
Beta (C2)	Financial Services	> 5.000	Germany	Normal	Insourced	Implemented	●	●	
Gamma (C3)	Infrastructure	> 4.000	Germany / Belgium	High	Insourced	Implemented	●	●	
Delta (C4)	Media	> 20.000	Germany	High	Hybrid	Interested	●	●	
Epsilon (C5)	e-Commerce	> 35.000	Germany	Normal	Hybrid	Intended	●		●

Figure 1. Case Data Overview.

We collected data over a ten-month period. Following Käss et al. (2024), we used multiple sources and data collection methods for triangulation. Specifically, we conducted nine interviews with key informants from AIPM, Data Science and Business Domain in four case organizations (Figure 1). The interviews lasted between 60 and 90 minutes each. All interviews were conducted via Microsoft Teams, recorded, and manually transcribed. In one case organization, we collected data by observing multiple internal meetings and a dedicated AIPM workshop (Figure 1). In addition, we analyzed documents from all case organizations. We anonymized the names of individuals and organizations before data analysis and presentation, including in this paper. Previous case studies on governing and managing AI have focused on small sample sizes to develop a deep contextual understanding. For example, Papagnannidis et al. (2023) conceptualize AI governance mechanisms by studying three organizations within the same industry. Therefore, we consider a sample of five organizations across industries sufficient for the purpose of our research.

3.3 Building the AI Portfolio Management Taxonomy

Taxonomies are built using two approaches: conceptual-to-empirical and empirical-to-conceptual (Nickerson et al., 2013). The former uses deductive reasoning to conceptualize characteristics and dimensions based on theory and creativity, then tests them against objects. The latter uses inductive reasoning to identify common characteristics among objects and groups them into dimensions. Both development processes iterate until the objective and subjective ending conditions are met. We grounded the development of our AIPM taxonomy in the data generated through the exploratory case studies. In an empirical-to-conceptual iteration, we used qualitative methods to identify common characteristics and group them into dimensions (Gioia et al., 2013). In a conceptual-to-empirical iteration, we applied

our (pre-)understanding of IT and other portfolio management disciplines. We conceptualized new characteristics and dimensions for the AIPM taxonomy and tested them against our data. We applied each development approach at least once (Kundisch et al., 2022). We built our taxonomy in four iterations (Figure 2) and summarize them as follows.

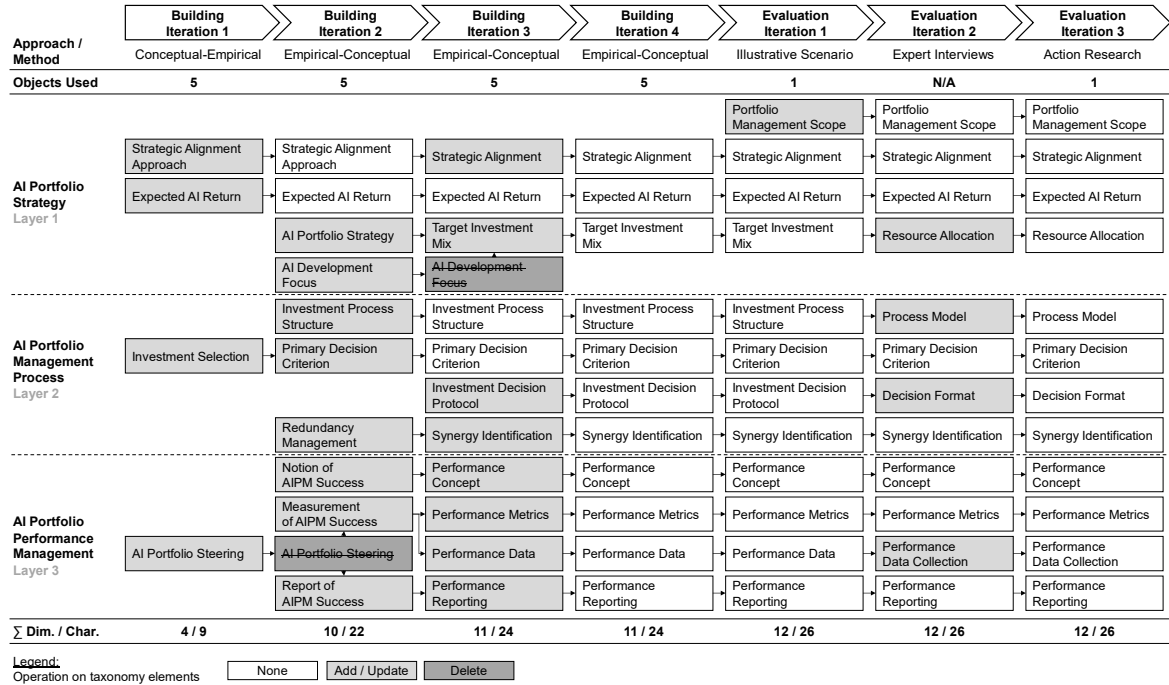


Figure 2. Evolution of our AIPM Taxonomy. Adapted from Schütz et al. (2023).

We chose the conceptual-to-empirical approach for the first iteration to build on existing knowledge from other portfolio management disciplines. For example, we drew on investment literature (Bierman, 1997) and assumed that organizations design AIPM for short- or long-term return. We created an initial AIPM taxonomy by testing these hypotheses against our data. We chose the empirical-to-conceptual approach for the second iteration. We identified seven new design dimensions by analyzing our case data inductively. The analysis also revealed new characteristics in existing dimensions. We continued with the empirical-to-conceptual approach in the third iteration. We identified one new, merged two and deleted one dimension. We also refined the wording of five dimensions and characteristics for clarity. We maintained the empirical-to-conceptual approach in the fourth iteration. This time, we did not identify new characteristics and dimensions or modify existing ones. The research team was confident in the taxonomy’s completeness, simplicity, understandability, and robustness. We agreed that both subjective and objective ending conditions were met and concluded the building process. We then moved to evaluate and refine the taxonomy.

3.4 Evaluating the AI Portfolio Management Taxonomy

We challenged the usefulness of our AIPM taxonomy through iterative evaluation. We used Szopinski et al.’s (2019) framework to configure appropriate strategies for each evaluation goal. For each evaluation, we varied the subjects, objects, and methods of evaluation (Szopinski et al., 2019). After each evaluation, we assessed whether the goals were met and refined our taxonomy as needed. The goals were met after three evaluations (Figure 3).

First, we assessed the analytical capabilities of our taxonomy through an illustrative scenario (Szopinski et al., 2019). We used the taxonomy to analyze an object that was not involved in building it. Specifically, AIPM at Swedish car maker Volvo as described by Andrén and Meddeb (2021). We found that our taxonomy lacked necessary elements. Volvo, unlike the organizations in our sample, uses portfolio management to manage AI initiatives from project proposal to deployment stage only (Andrén

& Meddeb, 2021). Consequently, we extended our taxonomy to account for different portfolio management scopes. The illustrative scenario also provided insights into how our taxonomy handles invalid inputs. Due to the secondary nature of the data, we could not classify Volvo across all dimensions and characteristics of our taxonomy. However, the research team agreed that acquiring more information would enable full classification. We thus considered our taxonomy robust. Because we added a new dimension and characteristics and did not meet all evaluation goals, we initiated another evaluation.

	Building Iteration 1	Building Iteration 2	Building Iteration 3	Building Iteration 4	Evaluation Iteration 1	Evaluation Iteration 2	Evaluation Iteration 3
	Conceptual-Empirical	Empirical-Conceptual	Empirical-Conceptual	Empirical-Conceptual	Illustrative Scenario	Expert Interviews	Action Research
Objective Ending Conditions							
All objects (or a representative sample) were analyzed			●	●	●	●	●
At least one object assigned to each characteristic	●	●	●	●	●	●	●
No dimensions / characteristics were added / modified in the last iteration				●			●
No dimensions / characteristics were merged / split in the last iteration	●			●	●	●	●
All dimensions and characteristic within each dimension are unique		●	●	●	●	●	●
Every combination of characteristics is unique			●	●	●	●	●
Subjective Ending Conditions / Evaluation Criteria							
Taxonomy contains all necessary elements / relationships (Completeness)			○	○		●	●
Taxonomy contains minimal number of elements / relationships (Simplicity)				○		●	●
Users comprehend the taxonomy and all its elements (Understandability)			○	○			●
Taxonomy handles invalid inputs / stressful conditions (Robustness)		○	○	○	●	●	●
Taxonomy Evaluation Strategy (Szopinski et al., 2019)							
Subject(s) of evaluation					Research Team Involved		
Subject(s): Involved in taxonomy building / Not involved in taxonomy building						2 / 6	0 / 3
Background: Academics / Practitioners					Academics	3 / 5	0 / 3
Experience: Domain / Method					Method	6 / 2	3 / 0
Object(s) of Evaluation					N = 1		N = 1
Type					Real-world	Real-World	Real-World
Object(s): Used / Not used in taxonomy building					Not used	Used	Not Used
Coverage					Representative	Representative	Representative
Method of Evaluation							
Approach					Qualitative	Qualitative	Qualitative
Method					Illustrative Scenario	Expert Interviews	Action Research

Legend: ● Ending Condition / Evaluation Goal Met ○ Ending Condition Assumed, Evaluation Required

Figure 3. Ending Conditions & Evaluation Goals Met in the Taxonomy Design Process.

Second, we conducted eight interviews with experts from academia and practice (Szopinski et al., 2019). Specifically, we interviewed three academics, two of whom had experience in taxonomy building and one who specialized in AI management. Additionally, we interviewed five practitioners who all had expertise in managing AI. We had interviewed two of these practitioners while collecting our primary data. Each interview lasted between 30 and 60 minutes. The experts reflected on the layers, dimensions, and characteristics of our taxonomy. They provided feedback on the terminology, whether the taxonomy included all necessary elements or if any could be removed. We also used the taxonomy to classify organizations within the design dimensions and characteristics with the two practitioners interviewed during data collection. Both experts confirmed that the taxonomy characterizes AIPM in their organization adequately. The evaluation indicated that our taxonomy contained all necessary elements without being overly detailed. We thus deemed the evaluation goals of completeness and simplicity as met. However, the experts suggested adapting the terminology of some characteristics and design dimensions for clarity. We thus revised the terminology and initiated another evaluation.

Third, we implemented an action research intervention (Szopinski et al., 2019). Specifically, we held a workshop with AI management and engineering specialists from a Swiss technology company. The practitioners had not participated in the development of the taxonomy. The goal of the full-day workshop was to reflect on the current management of the AI portfolio in the company and identify opportunities for improvement. We provided our taxonomy as a reflection tool. We introduced the design dimension and characteristics sequentially and illustrated them with examples from our case studies. Each practitioner independently classified the company according to the taxonomy, followed by a team discussion. After the workshop, we collected feedback on the terminology and assessed whether the taxonomy included all necessary elements or if any could be removed. The practitioners found the dimensions and characteristics understandable and the taxonomy easy to apply. No new characteristics

or design dimensions emerged during the workshop. Two months later, we held a reflection session with the company’s AI portfolio director. The director found the taxonomy useful as both an assessment tool and a conversation starter: “We created the greatest benefit [...] from this pigeon-hole representation of how we situate ourselves in the design dimensions [...]. After the workshop we were still talking about a few details where we disagreed.” The director also recognized the taxonomy as a discovery tool: “where I see value [...] is when you realize that there are design dimensions that you hadn't considered before.” The taxonomy contributed to the company’s strategic development: “Now we [...] have clearly delineated the roles of project manager, program director and also created a strategic role that is to take care of the design of structures and processes around AI portfolio management as such.” Based on the feedback, the research team considered the evaluation goals met.

Layer	Design Dimension	Characteristics			E/N
AI Portfolio Strategy	Portfolio Management Scope	AI Projects		AI Products	E
	Strategic Alignment	Business Unit Focus		Corporate Focus	N
	Expected AI Return	Short-Term Focus		Long-Term Focus	E
	Resource Allocation	Systematic		Ad Hoc	E
AI Portfolio Management Process	Process Model	High-Level		Low-Level	E
	Decision Format	Fixed		Flexible	E
	Primary Decision Criterion	Economic Viability	Technical Feasibility	Budget Availability	E
	Synergy Identification	Proactive		Reactive	E
AI Portfolio Performance Management	Performance Concept	Financial		Non-Financial	N
	Performance Metrics	Quantitative		Qualitative	N
	Performance Data Collection	Manual		(Semi)-Automated	N
	Performance Reporting	Retrospective	Prospective	No Formal Reporting	E

Legend : E = Exclusive Characteristic; N = Non-Exclusive Characteristic

Table 1. Multi-Layer Taxonomy of Design Decisions in AI Portfolio Management.

Based on the evidence from three evaluations, we have good reason to assume that our AIPM taxonomy fulfills its intended purpose. The final taxonomy consists of twelve design dimensions and 26 characteristics. The next section describes the taxonomy (Table 1) in more detail.

4 Taxonomy of Design Decisions in AI Portfolio Management

We identified decisions in twelve dimensions that organizations use to design portfolio management to govern AI. We grouped the design dimensions into three layers, without implying a hierarchy: *AI Portfolio Strategy*, *AI Portfolio Management Process*, and *AI Portfolio Performance Management*. We define all layers, dimensions, and characteristics in the following. We also provide examples from the organizations we studied. Our taxonomy is an analytical type of theory that “says what is” (Gregor, 2006). We thus use it to make sense of observed patterns, but not to explain why they occur.

4.1 AI Portfolio Strategy

AI Portfolio Strategy is the first layer in our AIPM taxonomy. It holds strategic decisions for the AI portfolio and its management to create business value. The layer contains four design dimensions: *Portfolio Management Scope*, *Strategic Alignment*, *Expected AI Return*, and *Resource Allocation*.

Portfolio Management Scope: Organizations design the portfolio management scope for AI projects or AI products. That means they apply portfolio management to a smaller or larger range of activities in the AI lifecycle. This depends on what AIPM is implemented for.

Organizations scope AIPM for *AI projects* to get the right systems into production. AI’s ability to learn and act without human intervention offers many opportunities to create business value. However, organizations face limitations in human, technological, and data resources. Therefore, they use portfolio management to decide which AI projects to pursue and to allocate resources efficiently. Because an AI

project portfolio focuses on development, its scope does not include post-deployment activities (e.g., monitoring and maintenance). Carmaker Volvo exemplifies this approach. Volvo's "Advanced Analytics & AI" department designed portfolio management for AI projects from proposal to system deployment (Andr n & Meddeb, 2021).

Organizations scope AIPM for **AI products** to ensure their AI systems create business value over time. This requires the portfolio management scope to include ongoing activities that monitor and maintain AI systems after deployment. This is because many AI systems learn from data and experience and their performance may degrade over time. As one respondent problematized: *"I don't know how extensively others define portfolio management. But there are a lot of people who put their AI models into production [...] and think 'let it run, my work is done. As a data scientist, I have nothing more to deal with'. But that's where it actually starts! When you have put an AI model into production, you need to constantly check: What is the difference to my control group? How much uplift is there really? [...] Do I need to retrain? When should I switch a model off again? [...] In my opinion, this is still neglected in 80% of the companies!"* (C-1). Applying portfolio management to a larger range of activities allows organizations not only to address risks that could reduce the business value of AI systems after they have been deployed. It also helps them identify opportunities to create new business value by improving or scaling AI products: *"You can either have a roadmap of what you want to develop further [...] or you can say: No, the product is good, I'm going to take care of the retraining, maintain it, and make sure I won't run into new risks"* (C-2). Interestingly, there is no consensus amongst the organizations we studied on when portfolio management should start to engage in AI product development. Financial services provider Beta considers activities such as *"filling an idea funnel"* or offering an *"ideation format [...] on the borderline"* (C-2). Whereas infrastructure provider Gamma holds *"AI demystification sessions"* to be *"absolutely critical"* (C-3). These early-stage activities define the business problems that AI products will address. Therefore, the extent to which portfolio management is involved determines how effectively AI products address key business challenges and align with overall strategy.

Strategic Alignment: Organizations focus the strategic alignment of their AIPM on corporate or business unit objectives. This decision depends on how much they expect AI to impact value creation.

Organizations align their AI portfolio with **corporate strategy** when they expect AI's advanced autonomy and learning capabilities to have a major impact on value creation. As a result, they set strategic goals for the AI portfolio and use it to drive change across the organization. For example, financial solutions provider Beta defined *"a consolidated goal that we're supposed to pay into and achieve across all [AI] products"* (C-2). Similarly, insurer Alpha sets *"the goals top-down e.g. we want X new signed policies or Y more net profit and the CDOs [...] derive how to achieve them with [...] AI models"* (C-1). Organizations align their AI portfolio with **business unit objectives** when they expect AI to impact certain aspects of value creation. Consequently, they do not set strategic goals for the AI portfolio but use it as a business unit proxy. As infrastructure operator Gamma explained: *"we [...] see ourselves more as [...] enabling the business to reach their strategic goals. I wouldn't say at the moment that we have our own goals"* (C-3).

Expected AI Return: Organizations design their portfolio management to focus on short or long-term AI returns. This depends on when they expect AI initiatives to create business value.

Organizations focus on **short-term return** when they need their AI portfolio to contribute to immediate business success. We found this focus to be common in highly competitive, digitalized industries such as insurance and financial services. High competition and digitalization pressure organizations to exploit AI's advanced autonomy and learning capabilities. As insurer Alpha stated: *"AI applications need to make money. When a division builds a new data science department, they have a grace period. But in the end, they have to make profit"* (C-1). Organizations focus on **long-term return** when they do not expect their AI portfolio to contribute to immediate business success. Instead, they expect it to create the necessary conditions to exploit AI's potential in the future. This involves experimenting with AI, building know-how, and exposing the workforce to technologies that are often complex and difficult to understand. Broadcaster Delta, for example, believes that *"AI [...] holds a certain promise [...] and because we want to explore exactly that, we are [...] willing to accept a lower cost-benefit-ratio"* (C-4).

Similarly, infrastructure provider Gamma uses its AI portfolio to “*show the business what machine learning can do*” (C-3).

Resource Allocation: Organizations allocate resources systematically or ad hoc. That means they use a more or less methodical approach to distribute their limited resources among the components managed in the AI portfolio (AI projects / AI products). This depends on how clearly they can articulate the ideal structure of their AI portfolio.

Organizations use the ***systematic*** allocation approach when they can formulate which types of AI projects / products their portfolio should ideally contain and how to distribute resources among them. The main objective is to ensure that each type of AI project / product has access to appropriate resources. This is because AI initiatives compete for the same limited resources, but have different characteristics: Some are “quick wins” that offer immediate business value at low cost (e.g., personalizing email newsletters), while others are “moon shots” that may require extensive research, long-term investment, and involve greater uncertainty (e.g., developing virtual shopping assistants). The AI portfolio risks favoring certain initiatives (e.g., quick wins over moon shots) if it does not take these characteristics into account. Organizations therefore construct certain AI project / product types to ensure that similar initiatives compete with each other (e.g., quick wins against other quick wins). Insurer Alpha, for example, classifies AI initiatives as either strategic or profit-driven and has defined “*a pot of money for [...] projects that are strategic to the company*” (C-1). Similarly, Volvo considered “strategic buckets” to manage its AI portfolio in the future (Andrén & Meddeb, 2021). Organizations allocate resources ***ad hoc*** when they have not (yet) defined their ideal AI portfolio. As a result, the AI portfolio is balanced and strategically aligned by selecting the most promising opportunities as they arise. This approach offers more flexibility than systematic resource allocation. We have found it to be common when organizations are still exploring their AI strategy or do not face major resource constraints. Financial solutions provider Beta exemplifies the ad hoc approach: “*We build teams around specific issues that we believe are important to us because they fit with the business. [...] Then, at shorter intervals, we think about what the relevant topics are. And then we have to see if these topics convince us. If so, they are allocated a budget.*” (C-2)

4.2 AI Portfolio Management Process

AI Portfolio Management Process is the second layer in our taxonomy. It contains decisions about the processes and practices that are put in place to realize the AI portfolio strategy. The layer includes four design dimensions: *Process Model*, *Decision Format*, *Primary Decision Criterion*, and *Synergy Identification*.

Process Model: Organizations design high-level or low-level process models for AIPM. That means they break the AIPM process into a smaller or larger number of activities. After each activity, they decide whether to continue, iterate, pause, or stop an AI project / product. The number of activities determines how often organizations (have to) make these decisions. Their preference for flexibility or control determines which process model they choose.

Organizations use ***high-level*** process models if they prefer flexibility in managing the AI portfolio. That is because AI’s ability to learn from data and experience makes many initiatives experimental. In early stages, it is often impossible to predict whether an AI initiative will achieve the intended goals and what the actual costs/benefits will be. In addition, some activities may need to be repeated (e.g. defining the business case, adjusting parameters, algorithms, or data inputs). High-level models reduce the number of activities and decision gates in the AIPM process. This creates flexibility and allows AI projects/products to iterate more freely and quickly. As a respondent explained: “[...] *we want to have as little discussion as possible about what status we are currently in. [...] That we get more into trying things out quickly. [...] That we think less in linear processes and more in iterations. [...] You often see portfolio drawings with four arrows. But actually, there’s always a back arrow involved*” (C-2). Financial solutions provider Beta, for example, uses two decision gates to manage all activities from ideation to decommissioning of AI systems: “[...] *we divided our portfolio into explore and exploit. That’s all we want, because we used to have more types, which made it harder than easier*” (C-2). For

this flexibility, organizations like Beta trade off some control over their AI portfolio as a whole. Conversely, organizations use **low-level** process models to gain more control and oversight of the AI portfolio. They achieve this by defining more activities and decision gates for AI initiatives to pass through in their lifecycle. Insurer Alpha, for example, uses three stages and decision gates to manage AI initiatives from ideation to deployment only: “*Proof-of-Concept, [...] Pilot [...] and Scale-Up*” (C-1). More decision gates provide more management opportunities. They also enable a more detailed view of the AI portfolio (e.g. for monitoring how many AI initiatives are currently at which stage, analyzing cycle times and success rates). For this gain in control and oversight of the AI portfolio, organizations like Alpha trade off flexibility in individual AI initiatives.

Decision Format: Organizations use fixed or flexible decision formats in the AIPM process. That means they formalize decision-making to varying degrees, leaving more or less discretion to decision-makers. This depends on what organizations believe will lead to good AIPM decisions.

Organizations use **fixed** formats to ensure consistent and transparent AIPM decisions. That is because AI’s broad application opportunities increase the number of people involved in or affected by decisions about AI projects/products in the portfolio. Therefore, organizations specify which techniques, artifacts, and information will be used in the decision process. Insurer Alpha, for example, uses “*business canvas and [...] AI technology canvas*” to inform “*a voting process where we also determine how confident we were about single metrics*” (C-1). Similarly, Volvo uses a set of weighted technical and business scores (Andrén & Meddeb, 2021, p. 43). These formalizations reduce the discretion of decision-makers. The general idea is that less discretion leads to decision outcomes that are more transparent and less dependent on individual judgment. We observed that formalization is particularly relevant when AIPM involves many decisions and decision-makers are expected to justify their choices. For example, when Volvo’s business units want the central AI department to explain which of the many AI project proposals it will (not) pursue and why. Organizations use **flexible** formats to adapt to the AIPM situation at hand. The idea is that by adapting the decision process, AIPM can capture important nuances of each decision situation and produce high-quality outcomes. This is achieved by giving decision-makers more discretion to use personal judgment and intuition. We observed that organizations tend to use flexible formats when AIPM decisions are infrequent, and the group of decision-makers is rather small. As a respondent from financial service provider Beta explained: “*I, as a product manager, pitch to my boss [...] and the better my arguments are, the more likely I am to succeed*” (C-2). What informs these decisions “*[...] is very specific. [...] The question is: Can I tell a good story about how I built my AI product and why I think it adds value? For that, I’ll take anything that’s useful.*” (C-2).

Primary Decision Criterion: Organizations use either economic viability, technological feasibility, or budget availability as primary criteria for AIPM decisions. That means they prioritize different aspects when deciding which AI projects/products in the portfolio to continue, iterate, pause, or stop. These decisions turn AI strategy into reality. Therefore, organizations choose the primary decision criterion to align with the strategic focus of the AI portfolio.

Organizations choose **economic viability** as the primary criterion when they expect the AI portfolio to contribute to immediate business success. They therefore allocate resources to AI projects/products that promise the best ratio of (potential) benefits to (expected) costs. Organizations try to express these costs and benefits in monetary terms whenever possible. This allows them to compare the business impact of AI projects/products that vary widely in technology and complexity. However, AI characteristics such as learning ability and data dependency create uncertainties that complicate the evaluation of costs and benefits. Especially in the early stages of AI innovation. As a respondent explained: “*[...] at the very beginning it’s unrealistic to estimate costs. [...] We try to estimate on a smaller scale and then make that decision over and over again*” (C-2). We found that organizations face similar challenges in determining business benefits. Specifically, in pinpointing how much business value AI projects/products generate on their own. There are two reasons for this. First, AI is often a supporting component within a larger IT system (e.g. product recommendations in an online store). Second, AI outputs (e.g. sales forecasts) can be used as input for many other business processes (e.g. procurement). Therefore, it may not be possible to quantify the business value created by individual AI projects/products. Despite these limitations, organizations choose economic viability as the primary

decision criterion to ensure that the AI portfolio addresses relevant business challenges and does not accumulate “[...] AI initiatives [that] are technically very interesting, but bring nothing at all in monetary terms” (C-1). While organizations like insurer Alpha and financial solutions provider Beta do consider other decision factors (e.g. technical feasibility), they view them as less important: “The first step is by benefit case: Can I say what the financial benefit is?” (C-2). Organizations use **technical feasibility** as the primary decision criterion when they do not expect their AI portfolio to contribute to immediate business success. Instead, they allocate resources to AI projects/products that are most likely to succeed given existing constraints (e.g. technology, data, infrastructure, expertise). Focusing on technical feasibility helps them to implement more AI systems and build competencies along the way. We found that technical feasibility is often a key driver in the early stages of AI adoption, where organizations focus on learning and have yet to detail a portfolio strategy. As broadcaster Delta explained: “At the very beginning, [...] we were very technology-driven [...]. We tried a lot of things and we were allowed to. [...] Now we evaluate the use case first and then select the technology. In the past, it was the other way around” (C-4). Organizations use **budget availability** as the primary decision criterion when the AI portfolio does not pursue its own strategic goals. Therefore, AIPM is not responsible for evaluating AI projects/products with regard to technical feasibility or economic viability. Instead, resources are allocated to AI projects/products based on the budget provided by other entities (e.g. business units). The AI portfolio is managed on the assumption that if and as long as a budget is provided, the AI project/product is valuable and should be pursued. As a respondent from infrastructure operator Gamma explained: “our idea is, [...] if the business comes with a budget, they have managed it against other things they could spend it on and decided that the value is enough for them. [...] They have to keep coming with a budget and if they don’t, we take that as a signal that it’s not giving enough value” (C-3).

Synergy Identification: Organizations take proactive or reactive approaches to identify synergies in their AI portfolio. This means that they use different strategies to find opportunities where different AI projects/products can support each other. Organizations design AIPM to detect potential synergies as they arise (reactive) or to seek them out before they become apparent (proactive). This decision depends on how much they focus on maximizing the business impact of the AI portfolio.

Organizations choose **proactive** synergy identification to increase the return on the resources invested in the AI portfolio. This means that they regularly review the portfolio for scenarios where different AI projects/products could benefit from each other. This is because developing and maintaining AI systems can be resource-intensive. We found that these organizations focus on value and cost synergies. Value synergies arise when AI projects/products can be scaled to solve additional business problems. Cost synergies occur when similar AI projects/products can be combined to save resources and reduce redundancy. We observed that organizations seek to identify these synergies proactively when they manage large AI portfolios and want to turn them into competitive advantage. Multinational insurer Alpha, for example, regularly screens the AI portfolio to identify scaling potentials: “Let’s assume I have discovered something in one country. The faster I can bring it to another country, the more competitive advantage I have” (C-1). Alpha also looks for cost synergies: “If you have a [...] portfolio with many similar use cases, it becomes super interesting [...]. If they use similar data and have similar outcomes, they can be consolidated!” (C-1). Managing large, diverse AI portfolios like Alpha’s requires “being able to exploit economies of scale.” (C-1) However, seeking out synergies requires additional effort which may not create enough value for other organizations. For them, it may be sufficient to react to synergies as they arise in the normal course of managing the AI portfolio. Organizations choose **reactive** synergy identification if they are less focused on increasing the return on the resources invested in their AI portfolio. Consequently, they design the AIPM process to increase the likelihood that synergies will emerge, but not to actively seek them out. We found that organizations take this approach when managing smaller AI portfolios in less competitive environments. They also tend to be more reserved about achieving cost synergies through consolidation. Infrastructure operator Gamma exemplifies this approach: “We tried to decrease redundancy. That is why we had these quarterly meetings [...] because we saw that they [the business units] are both often doing the exact same thing

separately and not talking to each other at all, which doesn't make sense. [...] However, what we don't want to do is force two different problems that are very similar to become one" (C-3).

4.3 AI Portfolio Performance Management

AI Portfolio Performance Management is the third layer in our taxonomy. It contains decisions about evaluating success and identifying opportunities to improve AIPM. The layer includes four design dimensions: *Performance Concept*, *Performance Metrics*, *Performance Data Collection*, and *Performance Reporting*.

Performance Concept: Organizations use financial or non-financial concepts to evaluate the performance of their AI portfolio. This depends on what they want to achieve by managing the AI portfolio.

Organizations use **financial** performance concepts when they manage their AI portfolio for monetary benefits (e.g. cost savings or profit generation). We found that this economic perspective is common in competitive industries with high pressure to adopt AI technologies. Financial performance concepts help organizations to invest in AI projects / products that support their financial health and strengthen their market position. Financial solutions provider Beta, for example, has defined "*a financial target [...] that we're supposed to [...] achieve across all [AI] products*" (C-2). Organizations adopt **non-financial** performance concepts when they do not expect their AI portfolio to contribute to immediate business success. Instead, they see AIPM as a way to build expertise and learn how to use AI technologies competently. We found that this broader understanding of performance is common among organizations that explore the impact of AI on their value creation and want to lay the groundwork to exploit economic potentials in the future. Infrastructure operator Gamma, for example, considers AIPM successful if it helps "*training [...] business product owners on how to manage data science products and spreading AI throughout the organization*" (C-3). Similarly, broadcaster Delta uses its AI portfolio for "*competence building [...] to make informed decisions about this whole [AI] matter. This discussion about ethics also plays a major role.*" (C-4).

Performance Metrics: Organizations use quantitative or qualitative metrics to assess the performance of their AI portfolio. The metrics depend on the performance aspects they want to evaluate.

Organizations use **quantitative** metrics to evaluate performance aspects that can be expressed through numerical indicators. Quantitative metrics allow them to compare individual AI projects / products and assess the development of the AI portfolio over time. Accordingly, we found organizations to favor quantitative metrics if they use AIPM to generate financial benefits (e.g. time and cost savings or profit gains). For example, insurer Alpha considers the aggregated net profit of the AI portfolio to be the "*most important KPI*" and also tracks average "*cycle times*" from ideation to deployment (C-1). Organizations use **qualitative** metrics to assess performance aspects that are difficult to quantify. These metrics provide contextual insights that help organizations evaluate the less tangible effects of managing the AI portfolio. Accordingly, we found that organizations tend to favor qualitative metrics if they use AIPM to achieve goals such as building AI competence in their workforce. Infrastructure operator Gamma, for example, uses qualitative feedback from business stakeholders: "*We have a survey that we send to the business every quarter for the different [AI] projects*" (C-3).

Performance Data Collection: Organizations use manual or (semi-)automated methods to collect data on the performance of their AI portfolio. This choice determines the amount of human effort involved and the type of performance information that can be gathered. Organizations thus select data collection methods that match how actively they manage their AI portfolio.

Organizations use **(semi-)automated** methods to collect performance data when they want to closely monitor their AI portfolio and make frequent adjustments. (Semi-)automation (e.g. by pulling data from webhooks and APIs) allows them to gather performance data at a scale and speed that manual approaches cannot match. As insurer Alpha explained: "*you can build stuff for portfolio steering relatively easily [...]. For example, a reporting API, [...] train a model [...] hook it into the API and the model reports continuously [...] how much improvement [...] it brought*" (C-2). This sort of (semi-)automation provides organizations with timely performance data that informs decisions about the AI

projects/products. We found that (semi-)automation is particularly useful for organizations with large AI portfolios that want to identify synergies early on. Like insurer Alpha: “*By not automating AI portfolio steering, you’re letting a lot of money slip through your fingers.*” (C-1). However, (semi-) automation requires upfront investments in time and resources, which may not be feasible for every organization. Organizations choose **manual** data collection when they do not need to closely manage the performance of their AI portfolio or make frequent adjustments. Manual data collection means that every piece of data about the AI portfolio is gathered by a human. Therefore, time and costs scale with the volume of data. We found that manual data collection is effective when organizations manage smaller AI portfolios, and timely performance information is less critical for decision-making. We also observed that the iterative nature of AI projects/products can make it difficult to keep information up to date. As financial services provider Beta explained: “*[A big challenge] is availability of data about what we are actually doing. If you’re a product manager, you have better things to do than always updating your business case*” (C-2). Organizations thus tend to reduce the effort of manual data collection by integrating it in established collaboration tools. Infrastructure Gamma, for example, has “*[...] a project pipeline in our Microsoft Teams channel [...]. And we have a Miro [online whiteboard] for the overview of forecasting and other initiatives*” (C-3).

Performance Reporting: Organizations generate retrospective, prospective or no formal reports on the performance of their AI portfolio. This choice depends on their strategic priorities and AIPM maturity. Organizations use **retrospective** reporting to evaluate the current business impact and operational success of their AI initiatives. Retrospective reporting involves regular reviews of the status and outcomes of projects/products within the AI portfolio. Insurer Alpha, for example, “*has a quarterly touchpoint with the main organizational units, where the units report on the projects they are currently working on and the goals [...] to be achieved*” (C-1). Retrospective reporting allows organizations like Alpha to compare the performance of the AI portfolio against strategic goals and identify opportunities to improve the AIPM process. Organizations use **prospective** reporting to strategize for future growth and innovation in the AI portfolio. Prospective reporting goes beyond describing the status quo and explores future potentials of AI initiatives. This involves techniques such as projections, simulations, and scenario analyses. Insurer Alpha, for instance, plans to “*link the KPIs to financial reportings [...]. Ideally, [...] we could start simulations in the AI portfolio to see how raising individual KPIs affects baseline financials or other metrics [...] to manage it through that as well*” (C-1). Similarly, financial service provider Beta foresees a future where AIPM answers: “*In which area could we go further to get the most out of our AI projects? [...] What could support each other?*” (C-2). Organizations have **no formal reporting** on the performance of their AI portfolio when they focus on institutionalizing basic AIPM strategies and processes. We found this to be true for most organizations in our sample as most of them had only recently begun to manage their AI projects / products as a portfolio. Financial solutions provider Beta and infrastructure operator Gamma, for example, noted that “*AI portfolio management itself does not yet generate reports*” (C-2) and considered formalized reporting “*an interesting idea*” (C-3). In the absence of formalized reporting, decision-makers rely on ad hoc updates to inform decisions about AI projects/products and continuously improve AIPM practices.

5 Discussion

We conducted this research to help understand how organizations design portfolio management to govern AI. We propose a multi-layered taxonomy to make sense of the key design decisions that organizations face in implementing AIPM. The taxonomy organizes these decisions into twelve dimensions across three layers: *AI Portfolio Management Strategy*, where organizations determine how the AI portfolio and its management will create business value. *AI Portfolio Management Process*, which establishes processes and practices to realize the strategy. And *AI Portfolio Performance Management*, which focuses on measuring success for continuous improvement. The taxonomy suggests that attention to all three layers is important for AIPM to be effective. We based our taxonomy on data from five cross-industry case studies. We evaluated and refined it over three cycles, incorporating feedback from eight expert interviews and a practitioner intervention. We propose our taxonomy as an

industry-agnostic resource to describe and analyze AIPM in organizations. Which we argue is useful for research and practice.

Theoretical Contributions: Our taxonomy of design decisions in AIPM makes three contributions to the body of knowledge on governance and the strategic use of AI in organizations (Birkstedt et al., 2023; Borges et al., 2021; Davenport, 2018; Fountaine et al., 2019; Haefner et al., 2023).

First, our taxonomy provides a lens to study organizational AIPM design both from specific and comprehensive perspectives. Previous research has tended to focus on specific AIPM issues, such as aligning AI initiatives with strategy (Borges et al., 2021; Davenport, 2018) or prioritizing them based on expected business impact (Davenport, 2018; Fountaine et al., 2019; Haefner et al., 2023). Our taxonomy recognizes that each of them involve important design decisions that can be studied independently. However, it also recognizes that these decisions are interdependent and should be studied as a whole. We argue that this dual perspective opens up two interesting avenues for advancing the current understanding of AIPM. On the one hand, our taxonomy encourages future research to explore how different design decisions influence one another. Effective AIPM, we suggest, requires coherent design decisions across *strategy*, *process*, and *performance management*. On the other hand, the dimensions of our taxonomy may help identify under-researched aspects of AIPM (e.g., KPIs for measuring the business impact of the AI portfolio). In both cases, our taxonomy can serve as a resource for developing new research questions or situating existing studies.

Second, our taxonomy provides a classification scheme that can help researchers study AIPM in organizations. We demonstrated its analytic capabilities in the results section, where we used it to describe how different organizations practice AIPM. Our taxonomy can support future descriptive and analytical research, as well as higher-level theorizing, in several ways. First, we observed that certain design decisions may co-occur (e.g., short-term return expectations and systematic synergy identification). Future research could use our taxonomy to identify clusters (Balijepally et al., 2011), archetypes, or successful configurations of AIPM. Second, researchers could examine tensions between dimensions of the taxonomy to explain why organizations make certain design decisions. Third, the taxonomy can be used to explore causal relationships between design decisions and the performance of an AI portfolio. Fourth, future research could expand our point-in-time perspective and investigate how organizations (re)-design AIPM over time, potentially leading to the development of maturity models. Ultimately, this cascade of research could lead to theories for design and action (e.g., methods, reference models; Gregor, 2006) that help academics and practitioners design and implement effective AIPM.

Third, our taxonomy contributes to the AI governance discourse by detailing the characteristics of portfolio management as an AI governance mechanism. It responds to recent calls for examining whether traditional IT governance mechanisms, like portfolio management, are suited for AI (Birkstedt et al., 2023). Our taxonomy shows that AIPM shares structural similarities with (IT) portfolio management with regards to *strategy*, *process*, and *performance management*. However, our research also highlights that organizations adapt certain aspects of (IT) portfolio management for AI. These adaptations are necessary because AI's ability to learn from data and experience introduces new uncertainties typically not encountered in (IT) portfolio management. Unlike most IT projects (e.g., deploying Microsoft 365), many AI initiatives face uncertainty regarding whether the intended outcomes can be achieved at all. This outcome uncertainty complicates portfolio management activities such as investment selection, strategic alignment, and resource allocation. Our taxonomy reflects these adaptations, for example detailing how organizations modify the process model to accommodate the exploratory and iterative nature of AI initiatives. Additionally, our research shows that organizations have multiple options for adapting portfolio management to AI-specific needs, such as defining high-level or low-level process models. As AIPM continues to evolve, we expect more real-world examples to emerge and therefore echo the call for ongoing research (Birkstedt et al., 2023).

Practical Implications: Our taxonomy also provides an actionable resource for practitioners, as evidenced by the intervention we undertook as part of our research. First, our taxonomy can act as a conversation starter. When we introduced it to practitioners, the taxonomy quickly prompted questions and constructive discussions, which helped to create a shared understanding. Second, practitioners can

use our taxonomy to evaluate their approach to AIPM by positioning it within the taxonomy's dimensions, discuss strengths and weaknesses, and identify areas for improvement. While engaging with practitioners, we have already observed discussions about key design decisions and the definition of measures or actions based on those discussions. Third, practitioners can use the taxonomy to compare their own AIPM approach to those of peer organizations. Therefore, we believe that our taxonomy can help more organizations than the ones used to evaluate its usefulness in building an effective AIPM.

Limitations: Our study has limitations that future research can address. First, we built our AIPM taxonomy based on five exploratory case studies. While appropriate for the purpose of our study, our sample may not fully capture the range of AIPM characteristics, potentially resulting in some being overrepresented and others underrepresented in our taxonomy. We therefore encourage future studies to employ representative sampling techniques and to evaluate our AIPM taxonomy in broader empirical settings. Second, we used qualitative methods for data analysis, which by nature introduce a certain degree of subjectivity. Future research could therefore diversify the data analysis methods. Recognizing the limitations of our work, we have designed our taxonomy to be adaptable, so others can revise it as the body of knowledge grows and AIPM evolves.

6 Conclusion

In this paper, we explored how organizations design portfolio management to govern AI. We contribute a taxonomy of design decisions that we propose as a tool to describe and analyze AIPM in organizations. We trust that researchers and practitioners will find our taxonomy useful in professionalizing AIPM. Which we believe is timely and important, as AIPM enables other governance activities. Such as managing AI risks, ensuring regulatory compliance, and that the use of AI contributes to overall success.

References

- Andrén, L., & Meddeb, J. (2021). *Project portfolio management for AI projects. Developing a framework to manage the challenges with AI portfolios*. Chalmers Open Digital Repository. <https://odr.chalmers.se/handle/20.500.12380/302494>
- Balijepally, V., Mangalaraj, G., & Iyengar, K. (2011). Are we wielding this hammer correctly? A reflective review of the application of cluster analysis in information systems research. *Journal of the Association for Information Systems*, 12(5), 375-413. <https://doi.org/10.17705/1jais.00266>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing Artificial Intelligence. *MIS Quarterly*, 45(3), 1433-1450. <https://doi.org/10.25300/MISQ/2021/16274>
- Bierman, H. (1997). Portfolio allocation and the investment horizon. *Journal of Portfolio Management: London*, 23(4), 51-55. <https://doi.org/10.3905/jpm.1997.409613>
- Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133-167. <https://doi.org/10.1108/INTR-01-2022-0042>
- Borges, A. F. S., Laurindo, F. J. B., Spínola, M. M., Gonçalves, R. F., & Mattos, C. A. (2021). The strategic use of artificial intelligence in the digital era: Systematic literature review and future research directions. *International Journal of Information Management*, 57, 102225. <https://doi.org/10.1016/j.ijinfomgt.2020.102225>
- Bosch, J., Olsson, H. H., & Crnkovic, I. (2021). Engineering AI Systems: A Research Agenda. In *Artificial Intelligence Paradigms for Smart Cyber-Physical Systems* (pp. 1-19). IGI Global. <https://doi.org/10.4018/978-1-7998-5101-1.ch001>
- Cheng, C.-C., Wei, C.-C., Chu, T.-J., & Lin, H.-H. (2022). AI Predicted Product Portfolio for Profit Maximization. *Applied Artificial Intelligence: AAI*, 36(1), 2083799. <https://doi.org/10.1080/08839514.2022.2083799>
- Chiang, I. R., & Nunez, M. A. (2013). Strategic alignment and value maximization for IT project portfolios. *Information Technology and Management*, 14(2), 143-157. <https://doi.org/10.1007/s10799-012-0126-9>

- Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.
- Fabozzi, F., Gupta, F., & Markowitz, H. (2002). The Legacy of Modern Portfolio Theory. *The Journal of Investing*, 11(3), 7-22. <https://doi.org/10.3905/joi.2002.319510>
- Fontaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-powered organization. *Harvard Business Review*, 97(4), 62-73. <https://hbr.org/2019/07/building-the-ai-powered-organization>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods*, 16(1), 15-31. <https://doi.org/10.1177/1094428112452151>
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611-642. <https://doi.org/10.2307/25148742>
- Haefner, N., Parida, V., Gassmann, O., & Wincent, J. (2023). Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197, 122878. <https://doi.org/10.1016/j.techfore.2023.122878>
- ISACA. (2019). COBIT® 2019 Framework: Introduction and Methodology. ISACA.
- Jeffery, M., & Leliveld, I. (2004). Best practices in IT portfolio management. *MIT Sloan Management Review*, 45(3), 41-49.
- Käss, S., Brosig, C., Westner, M., & Strahringer, S. (2024). Short and sweet: multiple mini case studies as a form of rigorous case study research. *Information Systems and E-Business Management*, 22(2), 351-384. <https://doi.org/10.1007/s10257-024-00674-2>
- Kundisch, D., Muntermann, J., Oberländer, A. M., Rau, D., Röglinger, M., Schoormann, T., & Szopinski, D. (2022). An Update for Taxonomy Designers. *Business & Information Systems Engineering*, 64(4), 421-439. <https://doi.org/10.1007/s12599-021-00723-x>
- Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1), 77-91. <https://doi.org/10.2307/2975974>
- Martínez-Plumed, F., Contreras-Ochando, L., Ferri, C., & Flach, P. (2021). CRISP-DM Twenty Years Later: From Data Mining Processes to Data Science Trajectories. *IEEE Transactions on Knowledge and Data Engineering*, 33(8), 3048-3061. <https://doi.org/10.1109/TKDE.2019.2962680>
- McKinsey & Company. (2023). *The economic potential of generative AI: The next productivity frontier*. Retrieved October 10, 2024 from <https://courses.cfte.education/digital-library-ai-mckinsey-report/>
- Nickerson, R. C., Varshney, U., & Muntermann, J. (2013). A method for taxonomy development and its application in information systems. *European Journal of Information Systems*, 22(3), 336-359. <https://doi.org/10.1057/ejis.2012.26>
- Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. *Information Systems Frontiers: A Journal of Research and Innovation*, 25(1), 123-141. <https://doi.org/10.1007/s10796-022-10251-y>
- PwC. (2024). *Thriving in an age of continuous reinvention*. Retrieved October 10, 2024 from <https://www.pwc.com/gx/en/ceo-survey/2024/download/27th-ceo-survey.pdf>
- Reyck, B. D., Grushka-Cockayne, Y., Lockett, M., Calderini, S. R., Moura, M., & Sloper, A. (2005). The impact of project portfolio management on information technology projects. *International Journal of Project Management*, 23(7), 524-537. <https://doi.org/10.1016/j.ijproman.2005.02.003>
- Schütz, F., Spierau, B., Rampold, F., Nickerson, R. C., & Trang, S. (2023). Chasing cyber security unicorns: A taxonomy-based analysis of cyber security start-ups' business models. Proceedings of the European Conference on Information Systems Kristiansand, Norway.
- Szopinski, D., Schoormann, T., & Kundisch, D. (2019). Because your taxonomy is worth it: Towards a framework for taxonomy evaluation. Proceedings of the European Conference on Information Systems Stockholm & Uppsala, Sweden.
- Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751-752. <https://doi.org/10.1126/science.aat5991>
- van Giffen, B., & Ludwig, H. (2023a). How Boards of Directors Govern Artificial Intelligence. *MIS Quarterly Executive*, 22(4), 251-272. <https://doi.org/10.17705/2msqe.00085>
- van Giffen, B., & Ludwig, H. (2023b). How Siemens Democratized Artificial Intelligence. *MIS Quarterly Executive*, 22(1), 1-21. <https://doi.org/10.17705/2msqe.00072>