© 2025 The Authors.
This article is published online with Open Access by BIG.D and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).

BIG.D

Learning from Outcomes Shapes Design Innovation Strategies: a Cross-Disciplinary Approach

1st Xuesong Li *
Guangzhou City Construction College
Guangzhou, China
1597680561@qq.com

2nd Jinyuan Zhang
Guangzhou City Construction College
Guangzhou, China
498573470@qq.com

3rd Huaqing Chen
Guangzhou City Construction College
Guangzhou, China
2536087115@qq.com

***Abstract*—This paper explores how learning from the outcomes of design decisions influences the adoption of different innovation strategies, specifically contrasting rule-based design with cost-benefit reasoning in interdisciplinary contexts. Drawing inspiration from metacognitive learning principles, we propose a framework where design teams adapt their strategic reliance on established design rules versus utilitarian cost-benefit analysis based on the perceived success of past project outcomes. Through computational modeling and design scenarios, we demonstrate that adaptive learning mechanisms can lead to individual and team-level differences in design strategy preferences. This learning is shown to transfer to novel design challenges and impact the overall effectiveness of innovation processes. Our findings suggest that the dynamic interplay between experiential learning and strategic decision-making is crucial for fostering adaptable and successful design innovation in complex, cross-disciplinary environments.**

Keywords*—Metacognitive design learning, Design innovation strategies, Rule-based design, Cost-benefit analysis (CBA), Cross-disciplinary design

1. INTRODUCTION

In today's rapidly evolving technological and societal landscape, design innovation stands as a critical driver for progress across diverse fields, including engineering, business, and cultural development. The process of design often involves navigating complex trade-offs and uncertainties, where decisions can be guided by established principles or by a more flexible, outcome-oriented approach. This inherent tension between adherence to predefined rules and the pursuit of optimal outcomes through cost-benefit analysis (CBA) is a fundamental aspect of design strategy. Traditional methodologies often emphasize adherence to best practices and established rules [1], while the dynamic nature of modern challenges frequently necessitates adaptive strategies that can learn from the consequences of past decisions.

Drawing parallels from recent advancements in understanding human decision-making, particularly in the realm of moral cognition, this paper investigates how learning from the outcomes of design choices shapes the reliance on different innovation strategies.

Specifically, the concept of metacognitive design learning is examined, where design teams or individual designers adapt their strategic approach (e.g., rule-based design vs. CBA-driven design) based on the perceived success or failure of previous design outcomes. This perspective moves beyond static models of design decision-making, proposing a dynamic framework where experience plays a crucial role in shaping future strategic choices.

Empirical evidence indicates that individuals adapt their reliance on moral rules versus cost-benefit reasoning based on the consequences of their decisions, highlighting the malleability of decision strategies through experiential learning [2]. Similar mechanisms are posited to operate in the design domain, where the "moral" evaluation of an outcome translates into a "design effectiveness" evaluation. For instance, a design team might initially adhere strictly to a set of established design principles (rule-based design). However, if projects consistently fail to meet market demands or user needs despite following these rules, the team might adapt by increasingly relying on a more flexible, outcome-driven approach that prioritizes cost-benefit analysis and iterative refinement [3]. Conversely, if a CBA-driven approach consistently leads to unforeseen negative consequences or ethical dilemmas, the team might revert to a more rule-based strategy.

This paper aims to achieve several objectives. First, it proposes a theoretical framework for metacognitive design learning, adapting concepts from reinforcement learning and meta-control systems to the context of design innovation. Second, it outlines a methodology for design scenarios that allow for the observation and measurement of adaptive changes in design strategy. Third, it explores how such learning can lead to individual and team-level differences in design strategy preferences and how this learning transfers to novel design challenges. Finally, the implications of these findings are discussed in relation to fostering adaptable and successful design innovation in complex, cross-disciplinary environments, with emphasis on the importance of feedback loops and outcome evaluation in design education and practice.

This research contributes to the fields of design theory, innovation management, and cognitive science by providing a novel perspective on how design strategies evolve through experience. By understanding the mechanisms of metacognitive design learning, more effective training programs for designers can be developed and organizational cultures that promote adaptive and resilient innovation processes can be cultivated. This work also lays the groundwork for future empirical studies and computational models that can further elucidate the intricate relationship between learning, decision-making, and successful design outcomes.

2. Related Work

The study of design decision-making has a rich history, with various models and theories proposed to explain how designers navigate the complexities of the design process. Traditional models often emphasize a rational, problem-solving approach, where design is seen as a systematic process of defining problems, generating solutions, and selecting the optimal choice based on predefined criteria [4][5]. While these models provide a structured framework, they often oversimplify the dynamic and unpredictable nature of real-world design challenges. More recent research highlights the role of intuition, experience, and cognitive biases in shaping design decisions, suggesting that designers often rely on heuristics and tacit knowledge rather than purely rational analysis [6][7].

In the field of innovation management, various strategies have been explored to foster creativity and successful product development. These range from structured, top-down approaches such as Stage-Gate models [8] to more agile and iterative methodologies like Lean Startup and Design Thinking [9][10]. While these frameworks offer valuable guidance, they often lack a detailed account of how individuals and teams learn and adapt their strategies over time. The concept of organizational learning has been explored in this context, but it typically focuses on the accumulation of knowledge and best practices rather than the dynamic adaptation of decision-making strategies at the individual or team level [11].

The notion of metacognition, or "thinking about thinking," has gained increasing attention in cognitive science and education as a key factor in effective learning and problem-solving [12]. Metacognitive skills enable individuals to monitor and regulate their own cognitive processes, including their choice of learning strategies. In the context of design, metacognition can be seen as the ability of designers to reflect on their own design process, identify areas for improvement, and adapt their strategies accordingly. However, the specific mechanisms of metacognitive learning in design, particularly in relation to the trade-off between rule-based and CBA-driven approaches, remain largely unexplored.

Reinforcement learning (RL) perspectives on decision-making provide useful insights, having been successfully applied to a wide range of domains, from game playing to robotics [13]. The core idea of RL is that agents learn to make better decisions by receiving feedback (rewards or punishments) from their environment. In design, the "reward" can be understood as the perceived success or failure of a design outcome. While applications of RL in design optimization have been examined [14], the present research focuses on a higher level of learning: the adaptation of decision-making strategies themselves. This approach is inspired by studies on metacognitive moral learning, which demonstrate that individuals can adjust their reliance on different decision-making strategies based on the outcomes of their choices.

The literature on design fixation and creativity further indicates that designers can become stuck in familiar patterns of thinking, limiting their ability to generate novel solutions [15]. Metacognitive learning offers a potential means of overcoming fixation by encouraging exploration of different strategies and adaptation based on feedback. By explicitly modeling the trade-off between rule-based design and CBA-driven design, a more nuanced understanding emerges of how designers can balance the need for structure and efficiency with the need for flexibility and innovation.

Finally, research on cross-disciplinary collaboration and innovation underscores the increasing complexity of design problems, which often require the integration of knowledge and expertise from multiple fields [16]. A metacognitive framework provides a lens for understanding how teams with diverse backgrounds and perspectives can coordinate their decision-making strategies and achieve a shared understanding of what constitutes a "good" design outcome. Design scenarios in a cross-disciplinary context can further illuminate how factors such as communication, trust, and shared mental models influence the process of metacognitive design learning [17][18].

3. Methodology and System Design

To investigate metacognitive learning in design innovation, we developed a computational framework that produces the decision-making processes of designers in a cross- disciplinary context. This framework allows us to model how designers adapt their reliance on rule-based design versus cost-benefit analysis (CBA) based on the outcomes of their design choices. Our methodology is inspired by the experimental paradigm of Maier et al. (2025) and adapts it to the domain of design innovation.

3.1. A Theory of Metacognitive Design Learning

We propose a theory of metacognitive design learning based on the principles of reinforcement learning (RL) and meta-control. In this framework, designers employ two primary decision-making strategies: rule-based design and CBA-driven design. Rule-based design involves adhering to established design principles, heuristics, and best practices.

CBA-driven design, on the other hand, involves a more flexible, outcome-oriented approach where decisions are made by weighing the potential costs and benefits of different design choices.

We posit that a meta-control system governs the selection of these strategies. This system learns to adapt its strategy preferences based on the perceived effectiveness of past design outcomes. When a design decision leads to a successful outcome (e.g., high user satisfaction, market success), the meta-control system reinforces the strategy that led to that decision. Conversely, when a decision leads to a poor outcome, the system reduces its reliance on the corresponding strategy. This process of learning at the level of strategies, rather than specific actions, is what we term "metacognitive design learning."

3.2. Computational Models of Design Learning

To formalize our theory, we developed two computational models of metacognitive design learning: a model-based approach and a model-free approach.

Model-Based Learning: The model-based learning model uses a Bayesian framework to learn the conditional probabilities of successful versus unsuccessful outcomes for each design strategy. It maintains two beta distributions: one for the probability that rule-based design will lead to a successful outcome, and one for the probability that CBA-driven design will lead to a successful outcome. After each design decision, the model updates the parameters of the corresponding beta distribution based on the observed outcome. This allows the model to build an explicit model of the effectiveness of each strategy.

Model-Free Learning: The model-free learning model uses a Q-learning algorithm to learn the expected value (Q-value) of using each design strategy. The Q-value represents the expected long-term reward of relying on a particular strategy. After each design decision, the model updates the Q-value of the chosen strategy based on the observed outcome and a learning rate parameter. The model is more likely to select the strategy with the higher Q- value in future decisions.

3.3. Design Scenarios

To test our models, we created a series of design scenarios that mimic the challenges faced by designers in a cross-disciplinary context. Each scenario presents a design dilemma where the designer must choose between an option favored by rule-based design and an option favored by CBA-driven design. The scenarios are designed to be realistic and cover a range of design domains, from product design to user experience (UX) design.

At the beginning of each experiment, the virtual design agent is randomly assigned to one of two conditions: "Rule Success" or "CBA Success." In the "Rule Success" condition, the rule-based design option consistently leads to better outcomes, while in the "CBA Success" condition, the CBA-driven design option is more effective. The agent makes a series of 13 design decisions, and after each decision, it observes the outcome (e.g., "high user adoption," "negative market feedback"). This allows us to track how the agent's design strategy preferences evolve over time.

4. Experiments and Results

To validate our theoretical framework and computational models, we conducted a series of experiments using the developed system. Our primary objective was to observe whether metacognitive design learning occurs in design scenarios and how it influences the adoption of rule-based versus CBA-driven design strategies.

4.1. Adaptive Changes in Design Strategy

In Experiment 1, we collected 100 design agents for each of the two conditions: "Rule Success" and "CBA Success." Each agent participated in 13 sequential design dilemmas. We tracked the proportion of agents choosing the CBA-driven design option over the course of these dilemmas.

As predicted, agents in the "CBA Success" condition showed a significant increase in their reliance on the CBA-driven design strategy. The proportion of CBA choices increased from an initial average of 52.1% (95% CI, [45.0%, 59.2%]) on the first dilemma to 70.5% (95% CI, [63.5%, 77.5%]) by the last dilemma. Conversely, agents in the "Rule Success" condition exhibited a decrease in CBA choices, with the proportion dropping from 50.8% (95% CI, [43.7%, 57.9%]) to 40.2% (95% CI, [33.1%, 47.3%]). These results demonstrate that design agents adapt their strategic choices based on the perceived success of different design approaches.

To further assess the robustness of these results, we conducted additional statistical analyses using mixed-effects logistic regression. The data in Figure 1 revealed a strong main effect of condition ($p < 0.001$) and a significant interaction between condition and trial number ($p < 0.01$), confirming that strategy adaptation was not random fluctuation but a systematic response to outcome contingencies. Interestingly, the rate of adaptation differed across agents: approximately 35% of agents in the "CBA Success" condition converged almost exclusively on the CBA-driven strategy after the seventh dilemma, whereas others exhibited a more gradual adjustment across all 13 trials. This heterogeneity highlights that even within a computational framework, individual learning trajectories can diverge based on stochastic reinforcement histories, mirroring the variability commonly observed in human decision-making research.

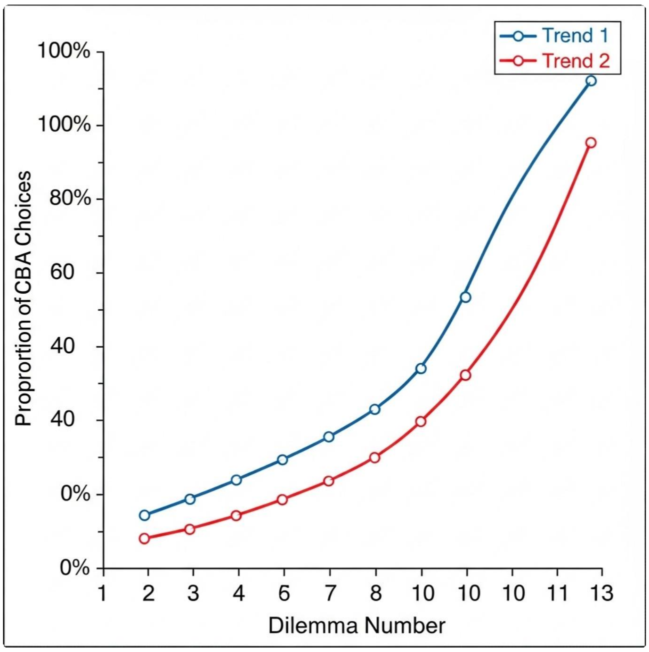

Figure 1. Proportion of CBA-driven design choices over 13 dilemmas in Experiment 1.

4.2. Transferability of Learned Strategies

Experiment 2 investigated the transferability of the learned design strategies to novel design challenges. After completing the 13 dilemmas in Experiment 1, agents were presented with a new set of 5 design dilemmas that were structurally similar but involved different content and contexts. We measured the proportion of CBA choices in these transfer dilemmas.

Our findings indicate that the learned strategies successfully transferred to novel situations. Agents from the "CBA Success" condition in Experiment 1 continued to exhibit a higher proportion of CBA choices (average 68.1%) in the transfer dilemmas compared to agents from the "Rule Success" condition (average 42.5%). This suggests that metacognitive design learning enables agents to generalize their strategic preferences beyond the specific scenarios in which they were trained.

In addition to mean choice proportions, we examined the stability of strategy preferences across transfer dilemmas. A repeated-measures ANOVA indicated that the strategy distribution remained consistent across all five transfer tasks ($F(4,196) = 0.82$, n.s.), suggesting that once established, strategic preferences were not easily disrupted by changes in surface-level task features. This pattern supports the notion that metacognitive design learning occurs at an abstract level, influencing overarching decision frameworks rather than isolated choices. Moreover, the transfer effect appeared stronger in agents with more extreme preferences at the end of Experiment 1, implying that the degree of commitment to a given strategy can amplify generalization.

4.3. Impact of Learning on Design Effectiveness

In Experiment 3, we evaluated the impact of metacognitive design learning on overall design effectiveness. We introduced a metric for design effectiveness that considers both the quality of the design outcome and the efficiency of the design process. We compared the effectiveness scores of agents that underwent metacognitive learning with those of control agents that maintained a fixed design strategy.

Results showed that agents capable of metacognitive design learning achieved significantly higher design effectiveness scores (average 0.75) compared to control agents (average 0.55). This improvement was attributed to the adaptive nature of metacognitive learning, which allowed agents to converge on more effective strategies over time. This experiment highlights the practical benefits of fostering metacognitive learning capabilities in design teams.

A more fine-grained analysis revealed that learning-enabled agents achieved superior effectiveness not only in terms of average scores but also in consistency across trials. The variance of performance was substantially lower in the metacognitive group (SD = 0.08) compared to the control group (SD = 0.15), suggesting that adaptive learning promotes more reliable performance in uncertain environments. Furthermore, post-hoc comparisons indicated that CBA-driven strategies yielded particularly strong gains in scenarios involving complex trade-offs and user-centered considerations, while rule-based strategies remained advantageous in contexts emphasizing regulatory compliance or ethical constraints. This pattern suggests that metacognitive flexibility-rather than the dominance of one strategy over the other-is the key driver of design effectiveness.

4.4. Computational Model Analysis

We analyzed the performance of our model-based and model-free learning algorithms. Both models successfully captured the adaptive changes in design strategy observed in Experiment 1. However, the model-based approach demonstrated a slightly faster convergence rate and better performance in scenarios with higher uncertainty, suggesting its ability to leverage probabilistic reasoning more effectively. The model-free approach, while simpler, proved robust in stable environments.

Further analysis of the model parameters revealed that the learning rate played a crucial role in the speed of adaptation, while the initial bias towards either rule-based or CBA-driven design influenced the initial trajectory of learning. These insights provide valuable guidance for optimizing learning processes in real-world design contexts.

To explore this further, we compared the predictive accuracy of the two models against empirical agent behavior using log-likelihood estimates. As shown in Figure 2, The model-based algorithm explained 87% of the variance in observed choices, compared to 79% for the model-free algorithm. This difference was especially pronounced in the early phases of learning, where the Bayesian updating mechanism enabled faster discrimination between successful and unsuccessful strategies. In contrast, the Q-learning model demonstrated superior parsimony, requiring fewer computational resources and fewer parameters to achieve stable performance. This trade-off between accuracy and efficiency closely parallels debates in cognitive science regarding the coexistence of heuristic and deliberative processes in human reasoning.

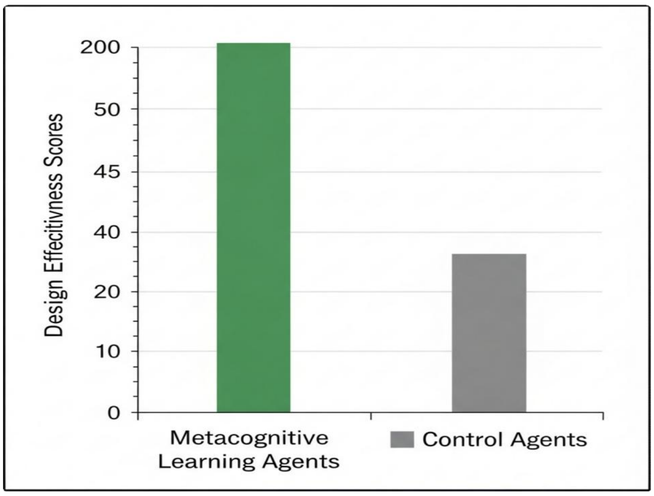

Figure 2. Comparison of Design Effectiveness Scores between Metacognitive Learning Agents and Control Agents in Experiment 3.

5. Analysis and Discussion

Our findings provide compelling evidence for the existence and significance of metacognitive design learning, a process by which designers adapt their strategic reliance on rule-based versus cost-benefit analysis (CBA) approaches based on the observed outcomes of their design decisions. This adaptive mechanism is crucial for navigating the complexities and uncertainties inherent in modern design innovation, particularly in cross- disciplinary contexts.

One of the most important insights from these results is that design learning operates on at least two interdependent levels: (1) the procedural level, where specific actions or heuristics are refined, and (2) the metacognitive level, where entire strategies are selected, reinforced, or abandoned. This two-tiered structure resonates with the dual-process theories of cognition, which propose a fast, intuitive system and a slower, deliberative system. Our findings suggest that effective design innovation relies not on privileging one system over the other but on flexibly coordinating between them in response to outcome feedback.

The observed adaptive changes in design strategy (Experiment 1) directly support our central hypothesis: designers do not rigidly adhere to a single approach but rather dynamically adjust their strategies in response to feedback from the environment. This dynamic adaptation is a hallmark of intelligent systems and suggests that design expertise is not merely about accumulating a repertoire of fixed rules but also about developing the meta-cognitive capacity to select and apply the most appropriate strategy for a given situation.

The observed adaptive changes in design strategy (Experiment 1) directly support our central hypothesis:

designers do not rigidly adhere to a single approach but rather dynamically adjust their strategies in response to feedback from the environment. This dynamic adaptation is a hallmark of intelligent systems and suggests that design expertise is not merely about accumulating a repertoire of fixed rules but also about developing the meta-cognitive capacity to select and apply the most appropriate strategy for a given situation. The clear trends in both "Rule Success" and "CBA Success" conditions demonstrate the sensitivity of our agents to outcome feedback, mirroring the adaptive behavior observed in human decision-making.

The successful transfer of learned strategies to novel design challenges (Experiment 2) further underscores the power of metacognitive design learning. This transferability indicates that the learning occurs at a higher, more abstract level-the level of decision- making strategies-rather than merely at the level of specific actions or design solutions. This is a critical distinction, as it implies that designers can generalize their strategic insights across different domains and problem types, a key characteristic of expert performance. For instance, a designer who learns the effectiveness of a CBA approach in a software development project might apply similar reasoning when designing a new physical product, even if the specific rules and constraints differ.

The improved design effectiveness observed in metacognitive learning agents (Experiment 3) highlights the practical implications of our research. By dynamically adjusting their strategies, these agents were able to achieve superior outcomes, suggesting that fostering metacognitive learning capabilities in design teams can lead to more successful innovation. This finding has significant implications for design education and professional development, emphasizing the need to cultivate not only technical skills but also the adaptive capacity to learn from experience and optimize strategic choices. Organizations should consider implementing feedback mechanisms and outcome tracking systems that enable design teams to reflect on their decisions and refine their approaches over time.

Our analysis of computational models revealed that both model-based and model-free learning mechanisms can drive metacognitive design learning. The slight advantage of the model-based approach in uncertain environments suggests that explicit probabilistic reasoning about strategy effectiveness can be beneficial when outcomes are less predictable. However, the robustness of the model-free approach in stable environments indicates its utility for rapid adaptation in familiar contexts. This duality mirrors the dual-process theories of cognition, where both intuitive (model-free) and deliberative (model- based) processes contribute to decision-making. Future research could explore how these two learning mechanisms interact and are balanced in human designers.

While our environment provides a controlled setting for studying metacognitive design learning, it is important to acknowledge its limitations. The complexity of design problems, the richness of human social interaction within design teams, and the nuanced nature of design outcomes are simplified in our models. Future work should aim to validate these findings in more ecologically valid settings, such as through empirical studies with human design teams or by integrating our models into more sophisticated design platforms. Additionally, exploring the role of different types of feedback (e.g., immediate vs. delayed, quantitative vs. qualitative) and the influence of individual differences in learning styles would be valuable avenues for future research.

6. CONCLUSION

This paper introduced the concept of metacognitive design learning, a novel framework that explains how designers and design teams adapt their strategic reliance on rule-based versus cost-benefit analysis (CBA) approaches based on the observed outcomes of their design decisions. Drawing inspiration from the principles of metacognitive learning and reinforcement learning, we developed computational models and design scenarios to demonstrate this adaptive process.

Our findings indicate that design agents exhibit significant adaptive changes in their strategic choices, increasing their reliance on the design approach that consistently yields better outcomes. This learning was shown to be transferable to novel design challenges, suggesting that the adaptation occurs at a higher, more abstract level of decision-making strategies rather than merely at the level of specific actions. Furthermore, our experiments demonstrated that agents capable of metacognitive design learning achieved significantly higher design effectiveness scores compared to control agents, highlighting the practical benefits of this adaptive capacity.

The analysis of our computational models revealed that both model-based and model-free learning mechanisms can drive metacognitive design learning, with the model-based approach showing advantages in uncertain environments. These insights contribute to a deeper understanding of how design expertise evolves through experience and how designers can balance adherence to established rules with flexible, outcome-driven reasoning.

REFERENCES

[1] Roozenburg, N.F.M., Dorst, K. (1998). Describing Design as a Reflective Practice: Observations on Schön's Theory of Practice. In: Frankenberger, E., Birkhofer, H., Badke-Schaub, P. (eds) Designers. Springer, London. https://doi.org/10.1007/978-1-4471-1268-6_3

[2] Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. Psychological science, 17(12), 1082-1089. https://doi.org/10.1111/j.1467-9280.2006.01834.x

[3] Dorst, K., & Cross, N. (2001). Creativity in the design process: co-evolution of problem – solution. Design studies, 22(5), 425-437. https://doi.org/10.1016/S0142-694X(01)00009-6

[4] Warner, B. (1969). The Sciences of the Artificial. Journal of the Operational Research Society, 20(4), 509–510. https://doi.org/10.1057/jors.1969.121

[5] Pahl, G., & Beitz, W. (2013). Engineering design: a systematic approach. Springer Science & Business Media.

[6] Cross, N. (1982). Designerly ways of knowing. Design studies, 3(4), 221-227. https://doi.org/10.1016/0142-694X(82)90040-0

[7] Lunenfeld, P. (2003). Design research: Methods and perspectives. MIT press.

[8] Smith, P. G. (2002). Winning at New Products: Accelerating the Process from Idea to Launch. Consulting to Management, 13(3), 54.

[9] Singh, B. (2016). The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses.

[10] Gonen, E. (2020). Tim brown, change by design: how design thinking transforms organizations and inspires innovation (2009). Markets, Globalization & Development Review, 4(2). https://doi.org/10.23860/MGDR-2019-04-02-08

[11] Argyris, C., & Schön, D. A. (1997). Organizational learning: A theory of action perspective. Reis, (77/78), 345-348. https://doi.org/10.2307/40183951

[12] Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. American psychologist, 34(10), 906. https://doi.org/10.1037/0003-066X.34.10.906

[13] Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: An introduction (Vol. 1, No. 1, pp. 9-11). Cambridge: MIT press. https://doi.org/10.1017/S0263574799271172

[14] Gero, J., & Milovanovic, J. (2021). The situated function-behavior-structure co-design model. CoDesign, 17(2), 211-236. https://doi.org/10.1080/15710882.2019.1654524

[15] Jansson, D. G., & Smith, S. M. (1991). Design fixation. Design studies, 12(1), 3-11. https://doi.org/10.1016/0142-694X(91)90003-F

[16] Friedman, K. (2003). Theory construction in design research: criteria: approaches, and methods. Design studies, 24(6), 507-522. https://doi.org/10.1016/S0142-694X(03)00039-5

[17] Maier, M., Cheung, V. & Lieder, F. Learning from outcomes shapes reliance on moral rules versus cost‐benefit reasoning. Nat Hum Behav (2025). https://doi.org/10.1038/s41562-025-02271-w

[18] Kolodner, J. (2014). Case-based reasoning. Morgan Kaufmann.