Human Factors, Business Management and Society, Vol. 215, 2026, 380–392
https://doi.org/10.54941/ahfe1007619

Efficient by Design: Using Behavioral Nudges to Eliminate Errors and Unlock High-Performance Operations

Joseph Caccitolo

The Chicago School of Professional Psychology, Chicago, IL, USA

Abstract

Organizations frequently respond to recurring errors, missed handoffs, delayed action, and rework as failures of individual attention, motivation, or accountability. Although accountability is necessary, many operational failures are better understood as predictable consequences of poorly designed decision environments. This conceptual and practice-oriented paper presents Efficient by Design, a human factors framework that applies behavioral economics, cognitive ergonomics, and choice architecture to everyday business management. The framework shifts leaders from a blame-oriented managerial stance to a diagnostic stance: first reconstructing how the work system made an error likely, then redesigning the environment so reliable action is easier. A central Decision Diagnostic is introduced to identify the decision being made, the user or worker making it, the cues and pressures present, the timing and quality of feedback, the easiest path, and the predictable error invited by the current system. Six redesign tools - defaults, feedback, mapping, error expectancy, incentives, and structured complex choice - are then linked to two human factors design goals: supporting fast System 1 thinking when routine action should be easy, and disrupting System 1 thinking when reflection, escalation, or expert judgment is required. Practitioner case illustrations from service delivery, technology selection, AI-enabled automation, and operational governance show how the framework can be applied to improve reliability, reduce cognitive load, and support sustainable growth culture in digitally transforming organizations.

Keywords: Human factors, Behavioral economics, Choice architecture, Decision environments, Business management, Organizational design, AI transformation

INTRODUCTION

Human factors research has long emphasized that error is rarely only a property of an individual person. Performance emerges from the interaction among people, tools, tasks, information, incentives, organizational norms, time pressure, and feedback (Norman, 2013; Reason, 2000). Yet in many business settings the first interpretation of an operational failure remains personal: someone missed the detail, ignored the process, failed to follow through, or did not care enough. The practical response is equally familiar: more training, more reminders, more oversight, or a more forceful demand for accountability. These responses sometimes help, but they also risk treating symptoms as causes. When the same failure recurs across capable people, the

Received June 14, 2026; Revised June 29, 2026; Accepted July 7, 2026; Available online July 20, 2026

more useful question is not simply who failed, but what in the system made the failure likely.

This paper presents the premise of the forthcoming book, *Efficient by Design*, as a human factors framework for business management and society. It is written for leaders, service operators, organization development practitioners, and researchers interested in how behavioral insights can be translated into practical methods for reducing operational friction. The argument is straightforward: many business errors are decision-environment problems. Work systems often ask people to perform accurately while under cognitive load, relying on memory, vigilance, interpretation of jargon, delayed feedback, and incentives that may reward speed, volume, or local optimization over reliability (Simon, 1955; Sweller, 1988; Wickens, 2008). Under those conditions, even skilled and motivated workers make predictable mistakes.

The AHFE Human Factors, Business Management and Society track emphasizes sustainable growth culture, AI-related digital transformation, leadership, organizational resource management, and practical methods for reducing complexity in management and leadership. *Efficient by Design* contributes to that agenda by offering a design approach that respects human cognition while improving business performance. This orientation is consistent with human-centered design principles that begin with a comprehensive understanding of users, tasks, and environments (International Organization for Standardization, 2019). Rather than asking people to be more perfect, it asks managers to design better work environments: clearer defaults, faster feedback, better mapping between action and consequence, anticipated errors, aligned incentives, and structured complex choices.

Conceptual Foundation: Why Capable People Make Predictable Mistakes

The framework begins with a core assumption shared by human factors, cognitive psychology, and behavioral economics: people are adaptive but bounded decision makers. They rely on attention, memory, mental models, habits, heuristics, social cues, and immediate feedback to navigate complex work (Simon, 1955; Wickens, 2008). These mechanisms are not flaws. They allow people to act quickly and efficiently in environments where deliberate analysis of every step would be impossible. However, when the environment hides important information, rewards the wrong action, makes the desired action difficult, or fails to warn the user at the point of risk, these same mechanisms produce error.

Dual-process theories are useful for explaining this pattern. Fast, intuitive, and automatic processes are well suited to routine tasks, pattern recognition, and familiar environments. Slower, deliberative processes are needed for ambiguous, high-stakes, novel, or exception-heavy decisions (Kahneman, 2011; Stanovich & West, 2000). Many business systems blur the distinction. They present high-risk decisions in routine-looking forms, place important information in dense dashboards, delay feedback until monthly reviews, and require users to remember exceptions while under pressure. The result is an environment that invites System execution when System 2 reflection is required.

Prospect theory and related work on framing also show that people do not evaluate outcomes in purely objective terms. Choices are shaped by reference points, perceived gains and losses, salience, and how options are described (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981). In business management, this matters because dashboards, service queues, procurement forms, project status categories, and AI recommendations all frame attention. A red indicator may feel urgent even when strategically minor. A default vendor may feel like the safe choice because it is already selected. A delayed cost may be ignored because it is not mapped to the current action.

Choice architecture extends these insights into design. If behavior is partly shaped by the environment in which choices are made, then managers and consultants can redesign the environment to improve the probability of better action (Johnson et al., 2012; Thaler & Sunstein, 2021). The ethical challenge is to do so transparently, preserving autonomy and professional judgment while reducing unnecessary cognitive burden. *Efficient by Design* treats choice architecture not as manipulation, but as a disciplined method for making desired work easier, predictable errors harder, and complex consequences more visible.

From Manager to Detective to Architect

A recurring failure often begins with a managerial question: who was responsible? The question is not illegitimate. Organizations require accountability, role clarity, and consequences. However, when the question comes too early, it narrows attention around the person closest to the failure. The Manager-Detective-Architect shift is proposed as a practical reframing device.

The manager asks, "Who is responsible?" The detective asks, "What actually happened, and what made it likely?" The architect asks, "How do we redesign the environment so better action becomes easier next time?" The Detective Mindset, inspired by the author's father's work as a police detective, emphasizes disciplined reconstruction before judgment. In organizational practice, this means examining cues, pressures, defaults, timing, handoffs, mental models, and feedback before concluding that the failure was primarily motivational or disciplinary.

The Architect Mindset then translates diagnosis into redesign. The goal is not to remove responsibility from people, but to create systems in which responsibility can be exercised effectively. This distinction is important in business management. A well-designed workflow does not infantilize professionals; it makes relevant information visible, reduces avoidable ambiguity, prompts attention at predictable failure points, and allows skilled workers to apply judgment where it matters most. This emphasis also aligns with psychological safety research showing that learning from failure depends on environments where people can speak up, report problems, and examine mistakes without reflexive blame (Edmondson, 1999; Frazier et al., 2017).

The Decision Diagnostic

The Decision Diagnostic is the central analytic method in *Efficient by Design*. It is intended for use before adding training, policies, or new technology. It asks leaders to identify the decision moment embedded in a visible problem. Instead of diagnosing "poor follow-through" in general, the diagnostic asks: What specific decision is the person making? What are they choosing between? What do they see at the point of action? What is the easiest path? When do they receive feedback? What error is predictable? What intervention would change the decision environment?

The diagnostic can be applied to frontline work, service operations, AI-enabled processes, procurement, project management, compliance workflows, and leadership routines. Its value is that it moves the conversation from a vague performance label to an observable decision situation. It also connects human factors analysis to practical management intervention. Using the diagnostic inform managers on how to intervene, leveraging tools that either support system fast thinking, or disrupt it, to engage thoughtful, effortful system 2. Table 1 summarizes the diagnostic questions and their redesign implications.

Table 1: Decision diagnostic questions and redesign implications.

Diagnostic Question	Human Factors Concern	Redesign Implication
What decision is being made?	The visible problem may hide a smaller choice point.	Define the exact moment to redesign.
Who makes the decision, and under what pressure?	Attention, workload, fatigue, expertise, and role clarity shape performance.	Fit prompts, defaults, and escalation to the actual user.
What cues are visible at the point of action?	Users act on salient information, not all available information.	Make the right cues timely, visible, and interpretable.
What is the easiest path?	Default effort often predicts behavior.	Make the preferred action easier or the risky action harder.
When is feedback received?	Delayed feedback weakens learning and calibration.	Move feedback closer to the decision.
What error is predictable?	Recurring errors reveal system affordances and traps.	Add guardrails, checks, or structured pauses.
Which intervention best fits the diagnosis?	Different failures require different design responses.	Select defaults, feedback, mapping, error expectancy, incentives, or structured choice.

Six Choice-Architecture Tools for Business Management

Efficient by Design organizes redesign around six tools. These tools are not meant to replace deeper organizational change, nor are they universal solutions. They are practical mechanisms for changing the decision environment at the point where work occurs. Each is taken from established Choice Architecture literature.

Defaults make the preferred action the easiest action without eliminating choice. In service operations, a ticket category, client assignment, or escalation path can be prepopulated based on reliable contextual information. Good defaults reduce variation and memory load; poor defaults silently reproduce error. Research on default effects, including organ-donation consent systems, demonstrates how strongly defaults can shape behavior while formally preserving choice (Johnson & Goldstein, 2003).

Feedback makes performance information visible while action is still possible. Many business systems provide feedback too late, such as month-end reports, quarterly performance reviews, or after-the-fact error lists. Human factors design favors immediate, action-linked, interpretable signals: warnings in forms, visible SLA timers, confirmation checks, or dashboards that distinguish urgent risk from cosmetic status. This is closely related to situation awareness, which depends on perceiving relevant elements, comprehending their meaning, and projecting their future status (Endsley, 1995).

Mapping connects current choices to future consequences. In technology selection, for example, a feature matrix may not help users imagine what implementation, maintenance, training, adoption, or total cost will feel like. Mapping translates abstract attributes into lived operational consequences: hours required, handoffs created, failure modes introduced, or user groups affected. It is used to avoid focalism, which can give authority attention to mundane features over individually meaningful ones.

Error expectancy assumes predictable mistakes will occur and designs for them. Instead of hoping users remember an exception, the system can prompt attention, constrain invalid entries, require confirmation, or insert an engineered pause. Error expectancy relies on either making the wrong choice impossible through physical or cognitive design elements or engineered pauses that demand multiple cognitive breaks for validation of accuracy or completeness. Error expectancy a aligns closely with human reliability and safety management traditions: recurring failure is information about the system (Reason, 2000).

Incentives shape attention and effort. They are used to align different actors within a system to nudge decision-making in support of system goals. They change the mental accounting such that actors' self-interest aligns with system interest. If teams are rewarded for closing tickets quickly, they may underinvest in knowledge capture or root-cause analysis. If sales teams are rewarded for speed, implementation teams may inherit poorly specified commitments. Incentive redesign asks whether personal, team, and enterprise goals point in the same direction. Application follows a *Certain – Simple – Immediate* format to ensure maximum salience and action.

Structured complex choice helps users make better decisions when options have many dimensions. Rather than asking a team to weigh every variable at once, leaders can specify criteria, cut scores, elimination rules, comparison matrices, and staged decision paths. This is especially relevant in AI-enabled business transformation, where technical possibility, risk, governance, cost, user trust, and workflow fit must be evaluated together.

Supporting or Disrupting System 1 Thinking

A key design decision is whether the system should support fast thinking or interrupt it. In routine, low-risk, high-frequency decisions, the goal may be to support System 1: make the correct action automatic, reduce unnecessary choices, and allow skilled workers to move quickly. Examples include the use of defaults, intuitive and error tolerant design. Business uses include standardized templates, pre-selected options, automated enrollments, any-orientation interfaces, and natural action-producing designs.

In novel, ambiguous, high-risk, or exception-heavy situations, the goal is to disrupt System 1. Here the system should slow action, make consequences visible, require structured reflection, or trigger escalation. Examples include confirmation prompts for high-cost decisions, warnings when data patterns suggest an anomalous entry, or required review before deploying an automation that affects customers. The same tool can serve either purpose. A default can speed appropriate action, while a structured pause can interrupt inappropriate automaticity.

This distinction is important for responsible AI implementation. AI can reduce cognitive burden by surfacing patterns, drafting responses, triaging work, and recommending next actions. But if users over-trust the output, accept defaults uncritically, or lose situational awareness, automation can amplify error. Automation research has long documented risks of misuse, disuse, abuse, and automation surprises when system behavior, authority, and user understanding are misaligned (Parasuraman & Riley, 1997; Sarter & Woods, 1997). *Efficient by Design* therefore treats AI not merely as a technology intervention, but as a decision-environment intervention. Managers must ask how AI changes cues, defaults, incentives, feedback, accountability, and human-machine balance.

Practitioner Case Illustrations

The following practitioner case illustrations are not presented as controlled experiments. They show how the Decision Diagnostic can translate human factors principles into management interventions by identifying the decision being made, the context in which it occurs, the predictable error, and the redesign tool most likely to improve outcomes. The cases are ordered according to the six tools introduced above: defaults, feedback, mapping, error expectancy, incentives, and structured complex choice, and represent only a sample of the business cases and real-world examples presented in the book.

Case 1: Defaults - AI-Enabled Service Delivery and the Risk of Making Automation the Default.

A managed IT services provider implemented automation to accelerate help desk ticket processing. The intent was reasonable: reduce manual effort, shorten processing time, and help analysts move tickets more quickly through intake and routing. The initial design, however, treated the AI-supported path as the operational default before the decision environment had been sufficiently tested. Analysts could accept suggested categories, priorities, and assignments with minimal friction, and the system made it easier to move forward than to pause, verify, or correct the recommendation. This created a classic default-risk problem. When an automated suggestion appears preselected, users may treat it as a cue of correctness, especially under time pressure or high ticket volume. Human factors research on automation cautions that automated systems can be misused or overtrusted when users defer to machine recommendations without adequate calibration (Parasuraman & Riley, 1997; Sarter & Woods, 1997).

The Decision Diagnostic reframed the failure. The decision was not simply whether analysts were following a new process; it was whether the system had made an unvalidated recommendation the easiest action. The predictable error was automation acceptance without sufficient review. The beneficial redesign was to change the default structure rather than abandon AI. The corrected design used automation to prepopulate low-risk fields only when confidence thresholds were met, required visible review for ambiguous categories, and routed exceptions to a human judgment point. In this configuration, AI supported System 1 for routine, high-confidence decisions while disrupting System 1 when the decision carried routing, priority, or client-impact risk. The outcome was not merely faster processing; it was more reliable processing because the system made the right degree of human attention the default.

Case 2: Feedback - Service Delivery and SLA Visibility.

A service delivery team repeatedly missed first-touch expectations on moderate-priority tickets. The initial managerial explanation was inconsistent follow-through: analysts knew the service-level agreement, but some tickets still aged past the target. The Decision Diagnostic shifted attention from motivation to visibility. The decision moment occurred while analysts were scanning a queue with many tickets that looked similar on the surface. The easiest path was to work the newest or most familiar item, not necessarily the item closest to breaching. Feedback existed, but it arrived too late through reports or after-the-fact escalation. In human factors terms, the system provided insufficient situation awareness at the point of action (Endsley, 1995).

The redesign used feedback as an immediate, action-linked signal. Tickets were visually coded by SLA risk so analysts could see which work required attention before the failure occurred. The intervention did not ask employees to try harder or remember more; it changed the perceptual environment so the relevant cue was available when the choice was being made. This supported System 1 rather than disrupting it: analysts could still move quickly, but fast attention was now guided by the right signal. The beneficial outcome was a reduction in missed first-touch expectations and a more constructive management conversation. Instead of asking why an analyst ignored the SLA, leaders could ask whether the work system made urgency visible early enough for a reasonable person to respond.

Case 3: Mapping - Technology Selection and Making Future Consequences Visible.

Technology selection often looks like an information problem. Buyers compare features, vendor demonstrations, pricing tiers, implementation promises, and reference calls. Yet the Decision Diagnostic showed that the real decision was not which product looked best in a presentation, but which future operating reality the organization was choosing. The predictable error was focalism: stakeholders overweighted visible, salient attributes such as a polished interface, impressive demonstrations, or a long feature list while underweighting the day-to-day consequences of configuration, adoption, support, reporting, data quality, training, and governance. This is a mapping failure because the choice format did not help decision makers connect current selection criteria to future lived experience.

The redesign translated abstract product traits into operational consequences. Instead of comparing vendors only by features, the team mapped each option to implementation burden, required behavior change, support model, failure modes, training requirements, reporting quality, and total cost of ownership. The Decision Diagnostic identified which stakeholders would live with each consequence and when feedback would arrive. This made hidden tradeoffs visible before commitment. The beneficial outcome was a more disciplined selection conversation: stakeholders could still value features, but they now evaluated those features alongside the work they would create. The mapping tool helped the group choose for welfare and operating fit, not only for presentation appeal or short-term enthusiasm.

Case 4: Error Expectancy - Designing an Engineered Pause Into High-Risk Workflow.

Some errors are predictable enough that a responsible system should expect them. In one operational workflow, the failure pattern was not that employees lacked competence; it was that the work design allowed a costly action to proceed with too little reflection. The Decision Diagnostic identified a high-frequency decision made under pressure, with several similar options, limited feedback before submission, and a costly downstream correction if the wrong path was chosen. The predictable error was a slip or shortcut created by cognitive load and routine action. Human error research emphasizes that reliable systems anticipate such conditions rather than relying only on vigilance (Reason, 2000).

The redesign added an engineered pause at the point where the mistake could still be prevented. Rather than adding broad training or another policy reminder, the system inserted a focused prompt, constraint, or confirmation only when the decision met a risk threshold. This disrupted System 1 at the right moment. The intervention was intentionally narrow: too many warnings would create alert fatigue, but a targeted pause for a known failure mode helped users shift from automatic action to deliberate review. The beneficial outcome was a reduction in avoidable downstream correction and a clearer distinction between routine work that should remain fast and exception work that should be slowed, checked, or escalated.

Case 5: Incentives - Motivating Help Desk Performance Through the CSI/WFH Case.

Incentive problems often appear as motivation problems, but the Decision Diagnostic asks whether the incentive is visible, attainable, timely, and connected to the behavior the organization wants. In a managed service desk of approximately 3analysts, the existing incentive was a monthly $25 bonus paid on the next paycheck for meeting a combination of performance requirements: 40 or more tickets, a mean time to resolution under two days, and 95 percent update compliance. In practice, almost nobody earned it. The incentive was delayed, uncertain, and psychologically distant. It also competed with the daily pressure to move through the queue. The behavioral problem was not that analysts did not value rewards; it was that the reward structure did not shape attention at the point of work.

The redesign kept the same performance metrics but changed the reward architecture. Analysts who met the requirements earned Friday work-from-home eligibility the following week. The new incentive was immediate, concrete, visible to the team, and highly valued. Approximately 8-10 analysts per week qualified, and performance improved. The Decision Diagnostic helped show why: the decision was not a one-time choice to be motivated, but repeated daily choices about updating tickets, resolving work, and maintaining quality. The WFH reward made the desired pattern personally meaningful and temporally close. This case also illustrates why incentives must be used carefully. Incentives can crowd out intrinsic motivation or distort behavior if poorly designed (Deci et al., 1999; Gneezy et al., 2011). Here, the benefit came from aligning a valued reward with balanced service measures rather than rewarding speed alone.

Case 6: Structured Complex Choice - Hiring a Store Manager.

Hiring a store manager is a structured complex choice because the decision requires comparison across many dimensions: operating discipline, customer orientation, leadership style, scheduling judgment, financial acumen, conflict management, culture fit, and growth potential. Without structure, hiring teams often rely on the most vivid interview moment, similarity to a successful prior manager, or general likeability. The Decision Diagnostic identified the decision as a high-impact, low-frequency choice with delayed feedback and many competing criteria. The predictable error was overweighting salient impressions and underweighting job-critical attributes that were harder to observe in an unstructured interview.

The redesign used structured choice to make the decision more comparable and less dependent on intuition alone. The hiring team first defined must-have criteria and cut scores, then separated screening criteria from development criteria, then used consistent questions and work-sample scenarios tied to the actual store manager role. Candidates were compared against the same dimensions rather than debated globally. This did not eliminate judgment; it improved the conditions under which judgment was exercised. Structured choice supported System 2 by slowing the team before the final decision and reducing the cognitive burden of weighing every trait at once. The beneficial outcome was a clearer, more defensible hiring process that improved alignment among decision makers and reduced the risk that the most charismatic or familiar candidate would be mistaken for the best operating fit (Iyengar & Lepper, 2000; Payne et al., 1993).

Implications for Sustainable Growth Culture

The Human Factors, Business Management and Society track emphasizes sustainable growth culture and shared value. *Efficient by Design* supports this orientation in three ways. First, it reduces complexity by giving managers a repeatable diagnostic method for operational friction. Second, it supports human-centered digital transformation by focusing attention on how technology changes decisions, workload, feedback, and trust. Third, it preserves dignity by treating error as a clue about the system before treating it as a defect in the person. Evidence that choice architecture can support less knowledgeable decision makers further supports the equity-oriented value of making complex systems easier to navigate (Mrkva et al., 2021).

This does not mean that every error is system-caused or that accountability is unimportant. Rather, the framework helps leaders sequence their response. Investigate the decision environment first; then determine whether training, accountability, technology, workflow redesign, incentive alignment, or cultural intervention is most appropriate. This sequencing is especially important in fast-changing organizations where AI, automation, distributed work, and complex service ecosystems create new sources of cognitive burden and coordination risk.

Limitations and Future Research

This paper presents a conceptual framework supported by established literature and practitioner cases. It does not report a controlled empirical test of the Decision Diagnostic or the six-tool model. Future research should examine whether teams using the framework achieve measurable reductions in error, rework, cycle time, escalation, or employee cognitive burden relative to teams using conventional training or accountability interventions. Research could also compare which tools are most effective for different categories of decision breakdown, such as attention failure, delayed feedback, poor mapping, incentive misalignment, or automation overreliance.

A second research need concerns ethics and organizational power. Choice architecture changes behavior by changing the environment. In business settings, this can support autonomy and performance, but it can also be used to manipulate, intensify work, or hide managerial priorities. Future work should define ethical criteria for transparent, participatory, and human-centered use of behavioral design in organizations (International Organization for Standardization, 2019; Thaler & Sunstein, 2021).

Conclusion

Efficient by Design offers a human factors approach to business management that reframes recurring operational failures not as evidence of individual shortcomings, but as valuable diagnostic signals about the design of work systems. By shifting leaders from a traditional managerial mindset ("Who failed?") to a Detective Mindset ("What in the system made this failure likely?") and finally to an Architect Mindset ("How do we redesign the environment so better performance becomes the default?"), the framework equips organizations with both a diagnostic lens and a practical redesign toolkit.

At its core, the Decision Diagnostic and the six choice-architecture tools — Defaults, Feedback, Mapping, Error Expectancy, Incentives, and Structured Complex Choice — provide a repeatable method for identifying where cognitive load, poor feedback, misaligned incentives, invisible consequences, or unnecessary complexity invite predictable errors. Rather than asking people to be more vigilant, more motivated, or more careful, the framework asks managers to shape the decision environment so reliable, high-quality action becomes the path of least resistance. This approach respects human cognitive architecture by supporting fast, intuitive System thinking for routine operations while strategically disrupting it when reflection, escalation, or expert judgment is required.

The implications extend beyond immediate error reduction. In an era of rapid digital transformation and AI adoption, thoughtful behavioral systems design is increasingly urgent. AI systems do not operate in isolation; they become powerful new elements within existing decision environments. Without deliberate design, automation can amplify existing biases, create new forms of over-reliance, or produce opaque workflows that reduce rather than enhance human agency. *Efficient by Design* treats AI not merely as a technical upgrade, but as a sociotechnical intervention that must be architected with human factors principles at its center.

For the Human Factors in Management, Business, and Society track, this framework contributes to sustainable growth culture by showing how organizations can reduce unnecessary friction, cognitive burden, and waste while improving reliability, employee experience, and long-term adaptability. It bridges academic behavioral science and daily operational reality, offering leaders a practical language and set of tools for building systems that work with human nature rather than against it.

Ultimately, *Efficient by Design* advances a simple proposition: organizational performance is shaped less by individual willpower than by the environments in which people make decisions. By embracing this human-centered perspective, leaders can move beyond reactive problem-solving toward proactive system stewardship, creating organizations that are not only more efficient, but also more resilient and humane.

REFERENCES

Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. Psychological Bulletin, 125(6), 627–668. https://doi.org/10.1037/0033-2909.125.6.627

DellaVigna, S., & Linos, E. (2022). RCTs to scale: Comprehensive evidence from two nudge units. Econometrica, 90(1), 81-116. https://doi.org/10.3982/ECTA18709

Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. Administrative Science Quarterly, 44(2), 350-383. https://doi.org/10.2307/2666999

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. Human Factors, 37(1), 32-64. https://doi.org/10.1518/001872095779049543

Frazier, M. L., Fainshmidt, S., Klinger, R. L., Pezeshkan, A., & Vracheva, V. (2017). Psychological safety: A meta-analytic review and extension. Personnel Psychology, 70(1), 113-165. https://doi.org/10.1111/peps.12183

Gneezy, U., Meier, S., & Rey-Biel, P. (2011). When and why incentives (don't) work to modify behavior. Journal of Economic Perspectives, 25(4), 191-210. https://doi.org/10.1257/jep.25.4.191

International Organization for Standardization. (2019). Ergonomics of human-system interaction - Part 210: Human-centred design for interactive systems (ISO Standard No. 9241-210:2019).

Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? Journal of Personality and Social Psychology, 79(6), 995-1006. https://doi.org/10.1037/0022-3514.79.6.995

Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? Science, 302(5649), 1338-1339. https://doi.org/10.1126/science.1091721

Johnson, E. J., Shu, S. B., Dellaert, B. G. C., Fox, C., Goldstein, D. G., Haubl, G., Larrick, R. P., Payne, J. W., Peters, E., Schkade, D., Wansink, B., & Weber, E. U. (2012). Beyond nudges: Tools of a choice architecture. Marketing Letters, 23(2), 487-504. https://doi.org/10.1007/s11002-012-9186-1

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263-291. https://doi.org/10.2307/1914185

Maier, M., Bartos, F., Stanley, T. D., Shanks, D. R., Harris, A. J. L., & Wagenmakers, E.-J. (2022). No evidence for nudging after adjusting for publication bias. Proceedings of the National Academy of Sciences, 119(31), e2200300119. https://doi.org/10.1073/pnas.2200300119

Mrkva, K., Posner, N. A., Reeck, C., & Johnson, E. J. (2021). Do nudges reduce disparities? Choice architecture compensates for low consumer knowledge. Journal of Marketing, 85(4), 67-84. https://doi.org/10.1177/0022242921993186

Norman, D. A. (2013). The design of everyday things: Revised and expanded edition. Basic Books.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230-253. https://doi.org/10.1518/001872097778543886

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). The adaptive decision maker. Cambridge University Press.

Reason, J. (2000). Human error: Models and management. BMJ, 320(7237), 768-770. https://doi.org/10.1136/bmj.320.7237.768

Sarter, N. B., & Woods, D. D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. Human Factors, 39(4), 553-569. https://doi.org/10.1518/001872097778667997

Simon, H. A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics, 69(1), 99-118. https://doi.org/10.2307/1884852

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? Behavioral and Brain Sciences, 23(5), 645-665. https://doi.org/10.1017/S0140525X00003435

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12(2), 257-285. https://doi.org/10.1207/s15516709cog1202_4

Thaler, R. H., & Sunstein, C. R. (2021). Nudge: The final edition. Yale University Press.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science, 211(4481), 453-458. https://doi.org/10.1126/science.7455683

Wickens, C. D. (2008). Multiple resources and mental workload. Human Factors, 50(3), 449–455. https://doi.org/10.1518/001872008X288394