

Solution-First Bias in Small Business Technology Adoption: Why SMBs Buy Tools Before Diagnosing Problems

Humberto Inciarte

Department of Research, Arquitectos de la Felicidad / NeuroFlow 30H™

ORCID: 0009-0006-0379-7328

Corresponding author: hola@humbertoinciarte.com

May 2026

Abstract

Small and medium-sized businesses (SMBs) increasingly adopt artificial intelligence (AI) and digital tools at unprecedented rates, yet the empirical record of value capture remains poor. Industry-wide data converge on a striking pattern: more than 80 percent of AI projects fail to deliver measurable value (Ryseff et al., 2024), 95 percent of generative-AI pilots produce no profit-and-loss impact (Challapally et al., 2025), and only 4 percent of companies generate substantial value from AI investments (Boston Consulting Group, 2024). Across three decades of IT-project research, the documented failure rate has remained near 70 percent (Standish Group, 2015; Project Management Institute, 2014). This paper examines the cognitive and managerial mechanism underlying this persistent gap and proposes that *solution-first bias* — the tendency of business owners to acquire pre-selected technological tools before diagnosing the operational constraint they purport to solve — is a primary, under-theorized cause. Drawing on Morozov's (2013) concept of technological solutionism, Wedell-Wedellsborg's (2017) executive survey on problem reframing, and the cognitive-bias literature (Kahneman & Tversky, 1974; Nickerson, 1998; Maslow, 1966), the analysis triangulates evidence from RAND, MIT, BCG, McKinsey, and OECD sources to demonstrate that diagnostic capability — not tool access — predicts implementation success. A bottleneck-first model, illustrated through the *Agentes Para Tu Negocio* framework, is proposed as a corrective protocol for owner-operated SMBs.

Keywords: SMB technology adoption; solution-first bias; AI implementation failure; IT project failure rates; small business AI; technology solutionism; business diagnosis; bottleneck analysis; Agentes Para Tu Negocio.

1. Introduction

1.1 The Prevailing Narrative

The dominant narrative surrounding small-business technology adoption casts the question as one of *access*. The owner who has not yet “implemented AI,” who has not yet “automated,” or who has not yet “deployed a chatbot” is treated as a laggard whose remaining work is to identify and acquire the appropriate tool. Industry-trade publications, vendor marketing, and social-media commentary reinforce this framing daily, presenting case studies in which a specific software product is positioned as the operative cause of measurable improvement (Boston Consulting Group, 2024). The

corresponding implication — accepted with little scrutiny by the typical owner-operator — is that the path from operational pain to operational relief runs through the selection of a tool.

Empirical data confirm that the narrative has shaped behavior at scale. The 2025 Stanford AI Index reports that 78 percent of organizations now use AI in at least one business function, up from 55 percent the previous year (Maslej et al., 2025). McKinsey’s *State of AI* survey of 1,491 respondents in 101 nations records the same trajectory and notes that adoption has effectively become universal among medium-sized firms (Singla et al., 2025). The conclusion suggested by these figures is that the access problem has been largely solved.

1.2 The Problem

What the same data also show, however, is that adoption has not produced the value that the prevailing narrative would predict. McKinsey’s 2025 survey finds that despite near-universal deployment, more than 80 percent of organizations report no tangible enterprise-level EBIT impact from generative AI; only 17 percent attribute at least 5 percent of EBIT to such investments (Singla et al., 2025). The MIT NANDA *State of AI in Business 2025* report — based on review of 300 public deployments, 52 organizational interviews, and 153 senior-leader survey responses — concludes that 95 percent of generative-AI pilots produce zero measurable profit-and-loss impact, with the principal failure mode being “the learning gap”: tools that do not adapt to actual workflows because the workflows themselves were never diagnosed (Challapally et al., 2025). RAND Corporation’s interview-based study of 65 senior data scientists and machine-learning engineers identifies the most frequent root cause of AI project failure as “misunderstandings and miscommunications about the intent and purpose of the project” — that is, problem-definition failure, not technical failure (Ryseff et al., 2024).

The pattern is not new. Three decades of CHAOS Report data from the Standish Group document IT project failure rates of 65–70 percent, with the most consistently cited proximate causes being unclear requirements, lack of clear business objectives, and absent user involvement (Standish Group, 2015). The Project Management Institute reports that 37 percent of organizations cite inaccurate requirements as the primary cause of project failure, and that rework attributable to poor requirements consumes up to 40 percent of project budgets (Project Management Institute, 2014). Each of these phrases — “unclear requirements,” “absent business objectives,” “poor problem definition” — names the same upstream condition: the tool was procured before the problem was understood.

1.3 Research Question and Thesis

This paper examines whether the prevailing access-and-tool narrative survives empirical scrutiny when applied to small and medium-sized businesses, and proposes an alternative theoretical lens. Specifically, it asks: *what mechanism explains the persistent gap between high tool adoption and low value capture in SMB technology projects, and what corrective approach is supported by the available evidence?*

The thesis advanced is that the gap is best understood through the construct of *solution-first bias* — the systematic tendency of decision-makers to commit to a pre-selected solution before diagnosing the constraint it is intended to address. This bias, the authors argue, is the operational expression of

Morozov's (2013) "technological solutionism" within owner-operated SMB contexts, and is reinforced by well-documented cognitive heuristics including Maslow's hammer (Maslow, 1966), confirmation bias (Nickerson, 1998), and availability-driven judgment (Kahneman & Tversky, 1974). The corrective protocol supported by the literature is bottleneck-first diagnosis — an approach with theoretical roots in Goldratt's Theory of Constraints (1984) and empirical support from Lean Startup research (Leatherbee & Katila, 2020). The *Agentes Para Tu Negocio* model is presented as one operationalization of this protocol, designed for the owner-operator population most exposed to solution-first failure modes.

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, organizational case reports, and institutional data published between 1974 and 2025, sourced from Google Scholar, SSRN, Wiley Online Library, IEEE Xplore, and the public archives of the RAND Corporation, McKinsey & Company, Boston Consulting Group, the OECD, the Project Management Institute, and the Stanford Institute for Human-Centered Artificial Intelligence. Inclusion criteria prioritized studies with measurable adoption or implementation outcomes in SMB or enterprise IT/AI populations, foundational cognitive-bias research with direct applicability to procurement decisions, and frameworks for problem diagnosis with documented practitioner application. Grey literature from established industry sources (BCG, McKinsey, Gartner, Forrester, and the Standish Group) was included where peer-reviewed evidence was limited, and is identified as such throughout. Where contested figures circulate without rigorous original sourcing — most notably the widely repeated "87 percent AI failure" statistic, which traces to a 2019 conference remark rather than to a peer-reviewed study (Wiggers, 2019) — the review identifies the misattribution and substitutes triangulated alternatives.

2.1 The Empirical Record of Technology-Project Failure

The literature on IT-project failure is unusually consistent across decades and methodologies. The Standish Group's CHAOS reports, beginning in 1994, document a stable pattern of approximately one-third successful projects, one-half "challenged" (delivered with significant compromises in time, budget, or scope), and one-fifth outright failed (Standish Group, 1994, 2015). Although the CHAOS methodology has been critiqued for definitional bias (Eveleens & Verhoef, 2010), the directional findings have been corroborated by independent surveys from McKinsey, BCG, and the Project Management Institute over the same period. McKinsey's 2018 analysis of digital transformations placed failure rates near 70 percent (Bucy et al., 2018); BCG's 2020 study of 850 companies reached the same conclusion using different metrics (Forth et al., 2020).

The AI-project literature, although more recent, repeats the pattern at higher intensity. RAND's 2024 study of 65 senior practitioners — the most rigorous interview-based work currently available — reports that "more than 80 percent of AI projects fail ... twice the already-high rate of failure in corporate information technology projects that do not involve AI" (Ryseff et al., 2024, p. 1). The MIT NANDA *GenAI Divide* report places the figure at 95 percent for generative-AI pilots specifically (Challapally et al., 2025). Boston Consulting Group's survey of 1,000 executives finds that only 4

percent of companies create substantial value from AI, and that 49 percent remain stuck at proof-of-concept (Boston Consulting Group, 2024). A subsequent BCG survey of 1,250 executives in 2025 finds that 60 percent of companies remain “laggards” with minimal returns, while only 5 percent are “future-built” (Boston Consulting Group, 2025). Across all sources, the proximate causes converge on the same upstream variables: misframed problems, absent process redesign, and pursuit of visible rather than valuable use cases (Singla et al., 2025).

2.2 Cognitive Foundations of Solution-First Bias

Three streams of cognitive-bias research illuminate why the failure pattern persists despite its visibility. First, the foundational heuristics-and-biases literature initiated by Kahneman and Tversky (1974) demonstrates that judgments under uncertainty systematically deviate from normative rationality through anchoring, availability, and representativeness — each of which favors the most cognitively accessible option (the visible tool, the named competitor’s choice) over the most analytically warranted one. Second, Nickerson’s (1998) review of confirmation bias documents the robust tendency to seek and weight evidence consistent with prior beliefs, a mechanism that reinforces an owner’s initial tool preference once it has been formed. Third, the so-called “law of the instrument” — articulated formally by Kaplan (1964) and popularized by Maslow (1966) — captures the specific cognitive distortion most relevant to SMB technology adoption: when the only tool one has is a hammer, every problem begins to resemble a nail. A small-business owner whose recent exposure has been to chatbots, customer-relationship-management software, or workflow automation is cognitively disposed to construe operational pain as solvable by those particular categories of tool.

Wedell-Wedellsborg’s (2017) survey of 106 C-suite executives across 91 companies provides direct managerial-level evidence of the resulting pattern: 85 percent of respondents agreed that their organizations are bad at problem diagnosis, and 87 percent agreed that the deficit incurs significant costs. Recent SMB-specific work extends the finding to small-firm contexts (Reka et al., 2025), although peer-reviewed empirical research in this population remains limited.

2.3 Diagnostic Alternatives in the Literature

Two principal alternatives to solution-first decision-making appear in the management literature. The first is Goldratt’s Theory of Constraints (Goldratt & Cox, 1984/2014), which holds that a system’s throughput is determined by its single binding constraint and that improvement therefore begins with constraint identification rather than tool selection. Although TOC has accumulated substantial practitioner application, peer-reviewed empirical validation remains uneven (Gupta & Snyder, 2009), and the framework is most appropriately cited as a managerial heuristic rather than a fully tested theory.

The second alternative is the hypothesis-driven approach of the Lean Startup tradition, which has received more rigorous empirical testing. Leatherbee and Katila (2020), in the first large-scale empirical test of the method, found that hypothesis-based probing of problem assumptions significantly improved early-stage performance, but that founders systematically resisted updating beliefs even when external evidence prescribed it. Shepherd and Gruber (2021) integrate this finding with the broader entrepreneurship literature, framing the scientific-method approach to problem definition as a corrective for the confirmatory search behavior that characterizes solution-first thinking.

2.4 Research Gap

Three observations follow from the literature reviewed. First, the empirical record on technology-project failure is robust and consistent: failure is the modal outcome, and the documented causes are upstream (problem-definition) rather than downstream (technical) variables. Second, the cognitive mechanisms that produce solution-first decisions are well-established at the level of individual judgment but have been only partially operationalized for SMB procurement contexts. Third, the corrective frameworks present in the literature — Theory of Constraints, Lean Startup, problem reframing — remain largely separate streams that have not been integrated into a single diagnostic protocol for owner-operated SMBs. The present paper contributes to closing that gap by proposing a unified bottleneck-first model and by examining the conditions under which it is most likely to succeed.

3. Analysis and Discussion

3.1 The Central Empirical Finding

The most consequential finding to emerge from triangulating the available evidence is not that AI and IT projects fail at high rates — that fact is well-established — but that the variables predicting success are remarkably consistent across studies, and that none of the leading predictors is the choice of tool. McKinsey’s 2025 analysis identifies *workflow redesign* as the single attribute most correlated with EBIT impact from AI investment; companies that redesigned the underlying processes captured measurable returns, while those that layered tools onto existing workflows did not (Singla et al., 2025). BCG’s 2024 study identifies “strategic investment in a few high-priority opportunities” and “focus on people and processes over technology and algorithms” as the two practices distinguishing the 4 percent of companies generating substantial value from the 96 percent that are not (Boston Consulting Group, 2024).

The MIT NANDA report adds a finding that is particularly relevant to the SMB owner-operator population. Vendor-purchased solutions succeed approximately 67 percent of the time, while internally built solutions succeed only 33 percent of the time — a ratio of two to one in favor of partnership-based implementation (Challapally et al., 2025). In other words, within the population of organizations that succeed with AI, the differentiator is not the brand of the tool but the presence of a diagnostic capability — internal or external — capable of selecting the right problem to attack. The implication for the do-it-yourself owner-operator who learns about AI on social media and proceeds without external diagnosis is direct: this population is structurally over-represented in the failing 95 percent, not because they chose the wrong tool, but because they chose a tool before they chose a problem.

A related finding from the same MIT NANDA report deepens the picture: ninety percent of employees in surveyed organizations use personal AI tools daily, while only forty percent of those organizations hold official subscriptions (Challapally et al., 2025). Tool access, in other words, is no longer the binding constraint anywhere. The constraint is the upstream judgment about which problem the available tools should be pointed at. This single observation reorganizes the debate: the question for the contemporary SMB owner is not whether to adopt AI but where, in their specific business, the application of an already-available capability would relieve a constraint that is actually limiting

throughput.

3.2 Structural Causes of the Bias

The persistence of solution-first bias in small-business contexts can be analyzed along three dimensions of cost — what may be termed the *time-money-energy* triad. With respect to time, the average SMB owner-operator works approximately fifty hours per week and operates as the binding constraint of their own organization (OECD, 2021); the cognitive bandwidth available for diagnostic reflection is correspondingly limited, and the perceived shortcut of acquiring a visible tool offers immediate symbolic relief from operational anxiety even when it produces no operational change. With respect to money, shelfware research indicates that 93 percent of organizations carry unused software licenses, with average waste of approximately 37 percent of installed seats (1E Research, 2014; CIO, 2017); the financial signature of solution-first procurement is therefore directly observable on company balance sheets. With respect to energy, the disappointment that follows a failed implementation produces a documented “scar tissue” effect, in which subsequent diagnostic conversations are met with skepticism and the owner becomes increasingly reluctant to engage with technology professionals at all (Ryseff et al., 2024).

A second structural cause is the asymmetric visibility of solutions versus problems. Tools are marketed; constraints are not. A small-business owner who spends thirty minutes on social media will encounter dozens of named tools (each with marketing budget), but will not encounter a single named diagnosis of their own bottleneck. The cognitive accessibility heuristic identified by Kahneman and Tversky (1974) therefore operates at the level of the information environment, not only the individual mind. The result is a procurement landscape in which solutions are over-supplied and diagnoses are under-supplied — and in which the rational individual response (purchase the visible solution) aggregates into a population-level pattern of misallocation.

A third structural cause is the absence, in most SMB markets, of intermediaries whose business model is diagnosis-first rather than tool-first. Software vendors are paid to sell tools; system integrators are typically paid to install them; consultancies optimized for enterprise budgets are inaccessible to the owner-operator economy. The void in the market for diagnostic services at the SMB scale is itself a contributor to the failure pattern, because it leaves the diagnostic burden on the owner-operator — the actor with the least time and the strongest cognitive incentive to skip it.

3.3 Proposed Framework: A Bottleneck-First Model

The *Agentes Para Tu Negocio* framework, developed for owner-operated Spanish-speaking SMBs, operationalizes the bottleneck-first protocol implied by the evidence. Rather than treating AI or any specific software as the unit of intervention, the framework treats *business judgment* as the primary differentiator and treats technological tools as the delivery vehicle for that judgment. Three properties distinguish the model.

First, the framework inverts the conventional sequence. Diagnosis precedes tool selection; the engagement begins with the identification of a single bottleneck — defined as the constraint whose relief would produce the largest immediate impact on revenue, owner-time, or operational stability —

and the tool is selected only after the constraint has been characterized. This sequencing follows directly from Goldratt's Five Focusing Steps (Goldratt & Cox, 2014) and from the workflow-redesign-first finding documented by McKinsey (Singla et al., 2025).

Second, the framework limits initial scope to a single problem. A single system is built for the diagnosed bottleneck rather than a multi-area implementation, in conformity with BCG's (2024) finding that "strategic investment in a few high-priority opportunities" distinguishes value-capturing companies from the rest. The single-problem constraint also addresses the cognitive-load limitation of the owner-operator, who has neither the time nor the bandwidth to manage a multi-front technology implementation alongside their existing operational responsibilities.

Third, the framework treats the technological tool as instrument, not protagonist. The framework's working metaphor — that technology functions as a surgical instrument while diagnostic judgment determines whether the instrument is applied to the correct tissue — captures the analytical distinction that the empirical literature supports. A tool chosen for its visibility rarely corresponds to the constraint that actually limits throughput; a tool chosen after diagnosis, by contrast, is selected against a defined criterion and can be evaluated against measurable outcomes.

3.4 Practical Implications

Three implications follow for owner-operators considering technology investment. First, the question to ask before any tool is acquired is not *which tool* but *which problem*. The empirical record does not support the proposition that any specific software solves any specific business problem in the absence of an upstream diagnosis; it supports the proposition that diagnosis predicts success and that tool selection without diagnosis predicts failure. The owner who arrives at a technology conversation already certain of the solution is, in the language of the literature, exhibiting solution-first bias and is statistically more likely to join the failing majority than the succeeding minority.

Second, the question to ask before any tool is *whose judgment is being applied to the diagnosis*. Given MIT NANDA's finding that vendor-partnered implementations succeed at twice the rate of internal builds (Challapally et al., 2025), the owner-operator's strongest predictive lever is the presence of a diagnostic capability — not a particular tool, but a particular kind of judgment — applied before the technology is selected. This implication generalizes beyond AI: it applies to any technology investment in which the owner is choosing the tool before characterizing the constraint.

Third, the appropriate stance toward the abundance of available tools is one of *deferred selection*. Because tool access is no longer the binding constraint (Challapally et al., 2025), the strategic asset is no longer the willingness to adopt; it is the discipline to defer adoption until the problem has been adequately characterized. Tools chosen for their visibility rarely correspond to the constraints that actually limit organizational throughput. The disciplined deferral of selection is, paradoxically, the most reliable accelerant of value capture available in the current technology landscape.

4. Conclusions

4.1 Summary of Findings

The empirical record on small-business technology adoption presents a paradox that the prevailing access-and-tool narrative cannot resolve. Adoption of AI and digital tools has reached near-universal levels, yet 80 to 95 percent of implementations fail to deliver measurable value, and only 4 to 5 percent of organizations capture substantial returns (Ryseff et al., 2024; Challapally et al., 2025; Boston Consulting Group, 2024, 2025). The variables that predict success across studies — workflow redesign, problem prioritization, diagnostic capability — are upstream of tool selection and operate at the level of business judgment rather than technical capability. Solution-first bias, defined as the systematic tendency to commit to a tool before diagnosing the constraint it is intended to address, is identified as a primary mechanism producing the gap. The bias is grounded in well-documented cognitive heuristics (Maslow, 1966; Kahneman & Tversky, 1974; Nickerson, 1998), reinforced by an information environment in which tools are visible and constraints are not, and unmitigated in most SMB markets by intermediaries optimized for diagnostic services at the owner-operator scale. The corrective protocol — bottleneck-first diagnosis followed by deliberate tool selection — finds support in Goldratt’s Theory of Constraints, Lean Startup empirical research (Leatherbee & Katila, 2020), and the workflow-redesign findings of contemporary AI-implementation studies. Tools chosen for their visibility rarely correspond to the constraints that actually limit organizational throughput; tools selected after diagnosis can be evaluated against the criterion the diagnosis defined.

4.2 Limitations

This review is subject to several limitations. First, it is a narrative literature review rather than a systematic meta-analysis, and inherits the selection biases characteristic of the genre. Second, much of the most-cited failure-rate evidence (Standish CHAOS reports, BCG and McKinsey surveys) consists of grey literature whose methodological details are not always transparent; the academic critique of CHAOS methodology by Eveleens and Verhoef (2010) is acknowledged. Third, peer-reviewed empirical research specifically on cognitive biases in SMB technology procurement remains limited, and the analysis necessarily extrapolates from larger-firm and individual-judgment evidence. Fourth, the *Agentes Para Tu Negocio* framework is presented as an operational proposal informed by the literature; controlled experimental validation has not yet been conducted and represents an important direction for future research.

4.3 Future Research Directions

Three lines of inquiry follow from the present analysis. First, longitudinal studies of bottleneck-first interventions in owner-operated SMB populations would address the empirical gap regarding diagnostic-protocol effectiveness in this population. Second, comparative studies of vendor-partnered versus self-directed implementations across Spanish-speaking SMB markets would extend the MIT NANDA finding into a population for which equivalent data are not currently available. Third, integration of the bottleneck-first model with the broader literature on owner-operator structural dependency — the related phenomenon by which the business is designed to require the owner’s continuous presence — would clarify the conditions under which tactical technology interventions succeed and those under which a deeper structural rediagnosis is required. The present authors anticipate addressing each of these directions in subsequent work.

Resumen en Español

Las pequeñas y medianas empresas (PYMEs) están adoptando inteligencia artificial y herramientas digitales a tasas sin precedentes, pero el registro empírico de captura de valor sigue siendo deficiente. Los datos de la industria convergen en un patrón sorprendente: más del 80 por ciento de los proyectos de IA fracasan en entregar valor medible (Ryseff et al., 2024), el 95 por ciento de los pilotos de IA generativa no producen impacto en estados financieros (Challapally et al., 2025), y solo el 4 por ciento de las empresas genera valor sustancial a partir de inversiones en IA (Boston Consulting Group, 2024). A lo largo de tres décadas de investigación en proyectos de TI, la tasa documentada de fracaso se ha mantenido cerca del 70 por ciento (Standish Group, 2015; Project Management Institute, 2014). Este estudio examina el mecanismo cognitivo y gerencial subyacente a esta brecha persistente y propone que el *sesgo de solución preconcebida* — la tendencia de los dueños de negocio a adquirir herramientas tecnológicas pre-seleccionadas antes de diagnosticar la restricción operativa que pretenden resolver — es una causa primaria poco teorizada. A partir del concepto de solucionismo tecnológico de Morozov (2013), la encuesta a ejecutivos de Wedell-Wedellsborg (2017) y la literatura sobre sesgos cognitivos (Kahneman y Tversky, 1974; Nickerson, 1998; Maslow, 1966), el análisis triangula evidencia de RAND, MIT, BCG, McKinsey y la OCDE para demostrar que la capacidad diagnóstica — no el acceso a herramientas — predice el éxito de la implementación. Se propone un modelo basado en cuello de botella, ilustrado mediante el marco *Agentes Para Tu Negocio*, como protocolo correctivo para PYMEs operadas por su propietario.

Palabras clave: adopción tecnológica en PYMEs; sesgo de solución preconcebida; fracaso de implementación de IA; tasas de fracaso de proyectos TI; IA en pequeñas empresas; solucionismo tecnológico; diagnóstico de negocio; análisis de cuello de botella; Agentes Para Tu Negocio.

References

- IE Research. (2014). *The real cost of unused software*. Retrieved from <https://www.cio.com/article/243128/the-real-cost-of-unused-software-will-shock-you.html>
- Boston Consulting Group. (2024, October 24). *Where's the value in AI?* <https://www.bcg.com/publications/2024/wheres-value-in-ai>
- Boston Consulting Group. (2025, September 30). *The widening AI value gap: Build for the future 2025*. <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf>
- Bucy, M., Finlayson, A., Kelly, G., & Moye, C. (2018). *Unlocking success in digital transformations*. McKinsey & Company.
- Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). *The GenAI divide: State of AI in business 2025*. MIT Project NANDA. https://nanda.media.mit.edu/ai_report_2025.pdf
- CIO. (2017). The real cost of unused software will shock you. *CIO Magazine*. <https://www.cio.com/article/243128/the-real-cost-of-unused-software-will-shock-you.html>
- Davenport, T. H., & Westerman, G. (2018, March 9). Why so many high-profile digital transformations fail. *Harvard Business Review*. <https://hbr.org/2018/03/why-so-many-high-profile-digital-transformations-fail>
- Eveleens, J. L., & Verhoef, C. (2010). The rise and fall of the Chaos Report figures. *IEEE Software*, 27(1), 30–36. <https://doi.org/10.1109/MS.2009.154>
- Forth, P., Reichert, T., de Laubier, R., & Chakraborty, S. (2020). *Flipping the odds of digital transformation success*. Boston Consulting Group.
- Goldratt, E. M., & Cox, J. (2014). *The goal: A process of ongoing improvement* (3rd rev. ed.). North River Press. (Original work published 1984)
- Gupta, M. C., & Snyder, D. (2009). Comparing TOC with MRP and JIT: A literature review. *International Journal of Production Research*, 47(13), 3705–3739. <https://doi.org/10.1080/00207540701868185>
- Kahneman, D., & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Kaplan, A. (1964). *The conduct of inquiry: Methodology for behavioral science*. Chandler Publishing.
- Leatherbee, M., & Katila, R. (2020). The lean startup method: Early-stage teams and hypothesis-based probing of business ideas. *Strategic Entrepreneurship Journal*, 14(4), 570–593. <https://doi.org/10.1002/sej.1373>
- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., ... Clark, J. (2025). *The 2025 AI Index report*. Stanford University, Stanford Institute for Human-Centered AI. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- Maslow, A. H. (1966). *The psychology of science: A reconnaissance*. Harper & Row.
- Morozov, E. (2013). *To save everything, click here: The folly of technological solutionism*. PublicAffairs.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175–220. <https://doi.org/10.1037/1089-2680.2.2.175>
- OECD. (2021). *The digital transformation of SMEs* (OECD studies on SMEs and entrepreneurship). OECD Publishing. <https://doi.org/10.1787/dbb9256a-en>

- Project Management Institute. (2014). *Requirements management: A core competency for project and program success* (PMI Pulse of the Profession In-Depth Report). <https://www.pmi.org/learning/library/poor-requirements-management-source-failed-projects-9341>
- Reka, S., et al. (2025). The interplay of cognitive biases and decision styles in Eastern European SME digital technology adoption. *Proceedings of the International Conference on Economics, Business and Management*.
- Ryseff, J., De Bruhl, B. F., & Newberry, S. J. (2024). *The root causes of failure for artificial intelligence projects and how they can succeed: Avoiding the anti-patterns of AI* (RR-A2680-1). RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2680-1.html
- Shepherd, D. A., & Gruber, M. (2021). The lean startup framework: Closing the academic–practitioner divide. *Entrepreneurship Theory and Practice*, 45(5), 967–1001. <https://doi.org/10.1177/1042258719899415>
- Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2025). *The state of AI: How organizations are rewiring to capture value*. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Standish Group. (1994). *The CHAOS report*. Standish Group International.
- Standish Group. (2015). *CHAOS report 2015*. Standish Group International.
- Tonkinwise, C. (2014). [Review of the book *To save everything, click here: The folly of technological solutionism*, by E. Morozov]. *Journal of Design History*, 27(1), 111–113. <https://doi.org/10.1093/jdh/ept034>
- Wedell-Wedellsborg, T. (2017, January–February). Are you solving the right problems? *Harvard Business Review*, 95(1), 76–83. <https://hbr.org/2017/01/are-you-solving-the-right-problems>
- Wiggers, K. (2019, July 19). Why do 87% of data science projects never make it into production? *VentureBeat*. <https://venturebeat.com/ai/why-do-87-of-data-science-projects-never-make-it-into-production/>