

Tacit Expertise as Competitive Differentiation: Why Generic AI Implementations Fail in Domain-Specific Business Contexts

Author: Humberto Inciarte

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

ORCID: <https://orcid.org/0009-0006-0379-7328>

Date: May 2026

Corresponding Author: hola@humbertoinciarte.com

Abstract

Following the broad public availability of generative artificial intelligence (AI) tools, a managerial narrative has emerged within the small and medium-sized business (SMB) community: that AI is now a commodity input, that any provider delivers equivalent results using the same underlying foundation models, and that the rational decision criterion has accordingly collapsed to price. This paper examines whether the commoditization thesis withstands empirical scrutiny. Drawing on the foundational literature on tacit knowledge (Polanyi, 1966; Nonaka, 1994), the resource-based and knowledge-based theories of the firm (Penrose, 1959; Barney, 1991; Grant, 1996), the cognitive science of expert judgment (Ericsson, Krampe & Tesch-Römer, 1993; Kahneman & Klein, 2009), and the recent body of empirical work on enterprise AI implementation outcomes (Brynjolfsson, Li & Raymond, 2025; Dell'Acqua et al., 2023; RAND Corporation, 2024; MIT NANDA, 2025), the analysis finds that aggregate failure rates between 80% and 95% in generative AI deployments cluster systematically around projects lacking integrated domain expertise, while vendor-led implementations succeed at approximately twice the rate of internal builds. The empirical record reframes generative AI not as a commoditized capability but as a delivery layer whose value is determined by the tacit business judgment encoded into its configuration. The paper proposes the *Agentes Para Tu Negocio* framework — a bottleneck-first, expertise-led implementation model for owner-operated SMBs — as a theoretically grounded corrective and identifies directions for further empirical validation in Latin American SMB contexts.

Keywords: tacit knowledge, AI customization, competitive differentiation SMB, expertise as moat, AI commoditization, knowledge management, business expertise, *Agentes Para Tu Negocio*

1. Introduction

1.1 The Prevailing Narrative

Since the broad public availability of large language models in late 2022, a particular framing of artificial intelligence has crystallized within the small and medium-sized business (SMB) community. Because the same foundation models — including GPT, Claude, and Gemini — are accessible to any business at low or zero marginal cost, AI is increasingly described as a commodity input: undifferentiated, broadly available, and therefore subject to the economic logic that governs commodities generally. Under this framing, any provider offering AI-related services is offering essentially the same thing as any other, and the rational decision criterion reduces to price.

The narrative is reinforced by the surface uniformity of the AI services market itself. Search results for "AI consultant," "custom chatbot," or "business automation" return functionally indistinguishable descriptions across providers, employing a converged vocabulary of "personalized chatbots," "intelligent automation," and "tailored AI solutions." Faced with linguistic and functional homogeneity, decision-makers default to the only available differentiator, namely cost. The phenomenon is consistent with the broader literature on

choice under uncertainty: when alternatives cannot be meaningfully discriminated, price becomes the dominant signal, even in cases where the underlying products or services are not, in fact, equivalent.

1.2 The Problem

The commoditization framing carries a strong empirical prediction. If AI services are genuinely commoditized, success and failure rates should be approximately uniform across providers, with cost-adjusted outcomes converging toward a stable mean. The empirical record contradicts this prediction sharply. The RAND Corporation's 2024 analysis of AI project outcomes reports that more than 80% of AI projects fail — twice the rate of comparable non-AI information technology projects (Ryseff, De Bruhl & Newberry, 2024). The MIT NANDA initiative's 2025 *State of AI in Business* report finds that approximately 95% of enterprise generative AI pilots produce no measurable impact on profit and loss, despite an estimated \$30–40 billion in enterprise generative AI expenditure (Challapally et al., 2025). The McKinsey 2025 *State of AI* survey, despite documenting near-universal adoption rates between 78% and 88%, finds that more than 80% of organizations report no enterprise-level EBIT impact from generative AI (McKinsey & Company, 2025). The Boston Consulting Group's 2025 review finds that only 4–5% of companies achieve substantial AI value at scale, with approximately 60% generating no material value despite investment (Boston Consulting Group, 2025). Gartner's 2026 survey of 782 infrastructure-and-operations leaders reports that only 28% of AI use cases meet ROI expectations, with 20% failing outright (Gartner, 2026).

The dispersion is the central anomaly. A genuinely commoditized input does not produce 80–95% failure rates while simultaneously enabling a 4–5% subset to achieve substantial returns. That distribution is the empirical signature of differentiated implementation, not of a commodity.

1.3 Research Question and Thesis

This paper examines whether the AI-as-commodity narrative withstands empirical scrutiny, and proposes an alternative framework — the *Agentes Para Tu Negocio* model — that addresses what the failure data reveal as the structural root cause. The thesis advanced here is that generative AI occupies the position of a vehicle, not of cargo: its value to a firm depends not on access to the underlying models, which is universal, but on the tacit business judgment encoded into its configuration, deployment, and governance. The apparent commoditization at the model layer makes applied judgment — the layer at which configuration, integration, and decision architecture occur — the operative competitive surface for SMBs.

The remainder of the paper is structured as follows. Section 2 reviews the foundational literatures on tacit knowledge, the resource-based and knowledge-based theories of the firm, and the cognitive science of expert judgment. Section 3 examines the recent empirical evidence on AI implementation outcomes, derives an analytical synthesis, and develops the proposed framework. Section 4 concludes.

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, foundational theoretical works, and tier-one institutional reports published between 1959 and 2026, sourced from Google Scholar, SSRN, PubMed, *Strategic Management Journal*, *Quarterly Journal of Economics*, *Organization Science*, *American Psychologist*, *Science*, and policy databases of the OECD, RAND Corporation, McKinsey & Company, the Boston Consulting Group, and the MIT NANDA initiative. Inclusion criteria prioritized (a) foundational theoretical sources on tacit knowledge and the resource-based view of the firm, (b) peer-reviewed empirical studies of generative AI productivity effects in business contexts, and (c) cross-sectional industry surveys of AI implementation outcomes with documented methodology and sample sizes above one hundred organizations. Grey literature was included where peer-reviewed evidence was limited and where the issuing institution provided documented methodology.

2.1 The Tacit Dimension of Knowledge

The conceptual foundation for the present analysis rests on Michael Polanyi's (1966) *The Tacit Dimension*, in which Polanyi advances the thesis that "we can know more than we can tell." Knowledge is structured asymmetrically: a substantial portion of what skilled practitioners know is held in a form that cannot be fully verbalized, formalized, or transmitted by explicit instruction. Polanyi's classic illustration involves the recognition of a face — readily performed but inarticulable — and extends to scientific judgment, technical skill, and what he termed "indwelling," the embodied familiarity through which experts navigate complex domains.

Nonaka (1994) operationalized Polanyi's framework for organizational contexts, formulating a dynamic theory of knowledge creation in which competitive advantage arises from the systematic conversion between tacit and explicit forms of knowledge. The four-mode SECI model — Socialization, Externalization, Combination, Internalization — locates the firm's distinctive capability in its ability to surface tacit expertise (externalization), recombine it with other explicit knowledge (combination), and re-embed the resulting synthesis into individual practice (internalization). Nonaka and Takeuchi (1995) extended the analysis to argue that the documented competitive advantage of major Japanese manufacturing firms in the 1980s and 1990s could be attributed largely to superior tacit-knowledge management practices.

Critically for the present analysis, Nonaka's framework identifies pure "combination" — the recombination of explicit knowledge into new explicit knowledge without intermediate tacit involvement — as systematically insufficient for generating sustained novelty or competitive advantage. Generic large language models, trained predominantly on publicly available text, perform exactly this combinatorial operation: they recombine explicit, codified knowledge. The tacit dimension that Nonaka identifies as decisive remains structurally external to the model.

Davenport and Prusak (1998), synthesizing two decades of organizational practice across hundreds of case studies, reach a converging conclusion. They argue that successful knowledge management is at most one-third a technological problem; the remaining two-thirds reside in people, processes, and culture. The most valuable knowledge — experience, intuition, beliefs — is consistently the most difficult to manage and to transfer. Their finding pre-figures by approximately a quarter-century the converging contemporary evidence that approximately 70% of AI value capture lies in workflow redesign and human change management, with 20% in technology and data and only 10% in algorithms (Boston Consulting Group, 2025).

2.2 The Resource-Based and Knowledge-Based Theories of the Firm

Penrose's (1959) *Theory of the Growth of the Firm* established the conceptual lineage that would, in subsequent decades, become the resource-based view of strategy. Penrose conceptualized the firm not as a black-box production function but as a bundle of idiosyncratic productive resources, the value of which arises not from the resources in themselves but from the firm-specific services into which they are converted by managerial judgment. The "Penrose effect" — the constraint that firm-specific managerial absorptive capacity places on growth — anticipates contemporary observations about the limits of generic AI capability without accompanying contextual integration.

Subsequent strategy scholarship formalized Penrose's intuitions into the resource-based view, arguing that competitive advantage is best analyzed from the resource side rather than the product side. Barney (1991) consolidated the framework with the criteria that have become canonical: a resource generates sustained competitive advantage only when it is Valuable, Rare, Imperfectly imitable, and Non-substitutable — the VRIN conditions. Barney further specified the conditions under which resources become imperfectly imitable: causal ambiguity, social complexity, and path-dependent unique historical conditions. Tacit knowledge, by Polanyi's definition, satisfies all three.

Grant (1996) extended the framework explicitly into the knowledge dimension, formulating the knowledge-based theory of the firm. Grant argues that the firm exists fundamentally as an institution for integrating specialized knowledge; tacit knowledge transfer is costly and imperfect, and the firm's coordination mechanisms exist precisely to economize on those costs. Teece, Pisano, and Shuen (1997) further developed the dynamic capabilities framework, arguing that sustained advantage in turbulent technological environments depends on the firm's capacity to sense, seize, and reconfigure resources — capabilities that are inherently path-dependent and firm-specific.

Applied to the present question, the resource-based and knowledge-based views jointly produce a clear analytical prediction. A generic AI model accessible to any firm satisfies the V condition (it is valuable) but fails the R, I, and N conditions: it is not rare, it is perfectly imitable, and it is easily substitutable. Such a resource cannot, in principle, generate sustained competitive advantage. Conversely, an AI implementation that integrates a firm's idiosyncratic tacit knowledge — its operational patterns, customer history, decision heuristics, and accumulated judgment — can in principle satisfy all four VRIN conditions, because the integrated configuration is causally ambiguous, socially complex, and path-dependent in Barney's sense.

2.3 The Cognitive Science of Expertise

A parallel literature in cognitive psychology has documented the structure of expert performance. Ericsson, Krampe, and Tesch-Römer (1993), in *Psychological Review*, argued that expert performance arises not from raw experience but from sustained engagement with structured deliberate practice over a period typically spanning approximately a decade. The findings have undergone partial replication and revision in later high-powered re-examinations, which confirm a real but smaller effect size; Ericsson himself disputed the popularized "10,000-hour rule" associated with the original study as a misrepresentation of the underlying claim. The defensible position from this literature is that expertise of the form relevant to business judgment — pattern recognition, anomaly detection, rapid contextual diagnosis — is acquired through extended structured engagement and cannot be shortcut.

Klein's (1993) recognition-primed decision (RPD) model, developed through naturalistic studies of firefighters, intensive-care nurses, and military commanders, complements Ericsson's account on the performance side. Klein finds that experts in time-pressured environments do not enumerate options and apply formal decision criteria; they recognize patterns and mentally simulate the first plausible course of action. The recognition itself draws on a tacit repository of prior cases, a structure incompatible with the kind of explicit reasoning that can be transferred via written instruction.

Kahneman and Klein (2009), in their notable joint paper *Conditions for Intuitive Expertise: A Failure to Disagree*, specify the boundary conditions under which expert intuition is reliable: (a) the environment must be sufficiently regular to be predictable, and (b) the expert must have had prolonged opportunity to learn those regularities through feedback. Crucially, they find that subjective experience is not by itself a reliable indicator of judgment accuracy; expertise is environmentally conditioned. The implication for the present analysis is dual. First, business judgment is real and acquirable, but only through extended exposure to a structured-feedback domain. Second, the same conditions that make expertise reliable in humans — prolonged exposure to firm-specific feedback environments — are precisely the conditions that distinguish a long-tenured operator from a generic AI system operating on broad, undifferentiated data.

2.4 Research Gap

The literatures reviewed converge on a coherent theoretical position: tacit knowledge is foundational to expertise (Polanyi; Nonaka); idiosyncratic, causally-ambiguous resources are the basis of sustained competitive advantage (Penrose; Barney; Grant; Teece et al.); and expert judgment is acquired through extended engagement with structured-feedback environments (Ericsson et al.; Klein; Kahneman & Klein). What the existing literature does not directly address is the application of these foundational findings to the specific question of how small and medium-sized businesses should structure their adoption of generative AI to convert generic-model access into firm-specific competitive advantage. The empirical literature on generative AI productivity, examined in Section 3, provides the missing evidence; the present analysis attempts to integrate the two.

3. Analysis and Discussion

3.1 The Empirical Failure Distribution

The aggregate evidence on enterprise AI implementation, sampled across multiple independent sources between 2023 and 2026, exhibits a remarkably consistent pattern. The RAND Corporation's interview-based study of sixty-five experienced data scientists and engineers (Ryseff, De Bruhl & Newberry, 2024) finds an AI project failure rate above 80%, and identifies the most common root cause as misunderstanding or miscommunication of project purpose — a failure of business judgment, not of technology. The MIT NANDA initiative's 2025 *State of AI in Business* report (Challapally et al., 2025) finds that 95% of enterprise generative AI pilots produce no measurable P&L impact. The McKinsey 2025 *State of AI* survey reports that more than 80% of organizations show no material EBIT contribution from generative AI, while a small subset (approximately 6%) qualify as "AI high performers" deriving more than 5% of EBIT from AI (McKinsey & Company, 2025). The Boston Consulting Group (2025) finds that only 4–5% of companies achieve substantial AI value at scale. Gartner's April 2026 release reports that only 28% of AI use cases meet ROI expectations, with 20% failing outright (Gartner, 2026).

Three features of this distribution warrant emphasis. First, the failure rates are convergent across methodologies — interview-based, survey-based, and outcome-based studies independently reach figures in the 70–95% range. Second, the small minority of successful implementations is consistently characterized by features that are not technological: workflow redesign, human-in-the-loop validation, deep contextual integration, and dedicated human-machine collaboration roles. Third, the MIT NANDA study's most differentiating finding — that vendor-purchased AI implementations succeed at approximately 67% versus 33% for internal builds — directly contradicts the commoditization thesis. If AI services were genuinely commoditized, vendor-led and internal builds would not differ by a factor of two.

3.2 The Jagged Frontier and the Distribution of Expertise Effects

The most rigorous experimental evidence on generative AI in knowledge-intensive business contexts comes from Dell'Acqua et al. (2023), whose field experiment with 758 Boston Consulting Group consultants on eighteen realistic tasks reveals what the authors term a "jagged technological frontier." Inside the AI capability frontier, GPT-4-assisted consultants completed 12.2% more tasks, 25.1% faster, and at higher quality than the no-AI control group. Outside the frontier — on tasks for which the model lacked the relevant capabilities — GPT-4-assisted consultants were nineteen percentage points less likely to produce correct solutions than the no-AI control. The authors describe this as "miscalibrated trust": knowledge workers systematically over-rely on AI in precisely the contexts where its limitations are most consequential. The detection of which side of the frontier a given task lies on is itself a function of contextual judgment that the model cannot perform on its own behalf.

Brynjolfsson, Li, and Raymond's (2025) *Quarterly Journal of Economics* study of 5,179 customer-support agents in a Fortune 500 deployment finds an average productivity gain of 14%, distributed asymmetrically: 34% gains for novice and low-skilled workers, with minimal gains for experienced and high-skilled workers. The authors interpret the result as evidence that AI assistance "disseminates the potentially tacit knowledge of more able workers" to less-experienced ones. The finding is doubly informative for the present analysis. First, it confirms empirically that what AI delivers in business contexts is, in measurable terms, the codified tacit knowledge of expert workers — not the model's autonomous capability. Second, it identifies the high-skill expert worker, not the model, as the primary source of the productive content being distributed.

Peng et al.'s (2023) field experiment with GitHub Copilot, finding 55.8% faster completion of an HTTP-server task with the largest gains accruing to less-experienced developers, reproduces the same pattern in a software engineering context. Noy and Zhang's (2023) *Science* study of ChatGPT use among professional writers finds a 40% reduction in task time and an 18% increase in quality, again with the largest effects among low-ability workers. These studies converge on a structurally consistent finding: generative AI tools produce substantial gains within their capability frontier, and those gains accrue disproportionately to workers whose own tacit knowledge is least developed. The implication for competitive differentiation is straightforward. AI tools level the commodity layer of work — the layer of routine task execution — while leaving the differentiated, judgment-intensive layer largely untouched.

The analytical synthesis is the following. Aggregate failure rates of 80–95% in enterprise AI deployments (Section 3.1) cluster around projects in which generic AI capability is deployed without integrated domain expertise. Successful deployments cluster around configurations in which expert tacit knowledge has been operationalized into the system. The 14% average productivity gain reported by Brynjolfsson et al. (2025) is real, but it is the average of a population that includes both successful (expertise-integrated) and unsuccessful (generic) deployments. The performance dispersion within and across firms is the empirical signature of differentiation, not of commoditization. The decision-maker who selects an AI services provider on the basis of price under the assumption of equivalence is, in measurable terms, optimizing against the wrong variable.

3.3 Proposed Framework: The Agentes Para Tu Negocio Model

The Agentes Para Tu Negocio model is proposed as a specific application of the integrated theoretical framework — Polanyi and Nonaka on tacit knowledge, Barney and Grant on the resource and knowledge bases of competitive advantage, and Klein and Kahneman on the structure of expert judgment — to the population of owner-operated small and medium-sized businesses, with particular relevance to the under-served Spanish-speaking SMB market identified in OECD (2024) data. The OECD finds that across seven member countries, the dominant adoption barrier for SMBs is the absence of skills (cited by 50% of respondents) and the absence of contextual knowledge (cited by 40–80% by country) — not the absence of access to the technology itself.

The model is structured around four propositions, each grounded in the literatures reviewed above:

Proposition 1 (Bottleneck-First Diagnosis). AI implementation should be preceded by a structured diagnosis of the firm's most binding operational bottleneck — the point at which the firm's specific resources, in Penrose's (1959) sense, are most constrained — rather than by tool selection. This sequencing is consistent with the OECD (2024) finding that contextual knowledge gaps, not technology access, dominate SMB adoption barriers, and with Ryseff et al.'s (2024) finding that the modal cause of AI project failure is misunderstanding of project purpose.

Proposition 2 (Tacit-to-Explicit Codification). The implementation process should externalize, in Nonaka's (1994) sense, the operator's accumulated tacit business knowledge — pattern recognition concerning customer behavior, decision heuristics for pricing and operations, recognition-primed responses (Klein, 1993) to recurring situations — and encode that knowledge into the structural prompts, evaluators, and decision architectures of the AI system. The deliverable is therefore not a configured tool but a codified judgment artifact.

Proposition 3 (Frontier-Aware Deployment). Deployment must respect Dell'Acqua et al.'s (2023) jagged frontier: the AI system is deployed in contexts where the model's capability is established (high-volume, well-bounded tasks) and explicitly excluded from contexts beyond its capability frontier, with human judgment retained for ambiguous, high-consequence decisions. This is the operational expression of the Brynjolfsson et al. (2025) finding that AI's productive role is the dissemination of expert tacit knowledge to lower-skilled tasks, not the autonomous replacement of expertise itself.

Proposition 4 (Vendor-Led Integration). Consistent with the Challapally et al. (2025) finding that vendor-led implementations succeed at approximately twice the rate of internal builds, the model positions an external implementer not as a tool installer but as the carrier of integrating tacit knowledge into the firm's specific configuration. The external party brings business and operational judgment that the internal organization, by hypothesis, does not yet possess in encoded form; the implementation process is the structured transfer and codification of that judgment.

The four propositions jointly satisfy the VRIN conditions specified by Barney (1991). The resulting configuration is Valuable (it produces measurable operational improvement), Rare (the integrating expertise is itself scarce), Imperfectly imitable (the encoded tacit knowledge is causally ambiguous and path-dependent), and Non-substitutable (generic AI access does not produce equivalent firm-specific outcomes). The model thus offers a theoretically grounded route by which a small or medium-sized business can convert generic-model access — a non-VRIN resource — into a firm-specific configuration that does satisfy the conditions for sustained competitive advantage.

3.4 Practical Implications

Three practical implications follow from the foregoing analysis.

First, the price-based decision criterion that emerges from the commoditization narrative is empirically miscalibrated. When the dispersion of outcomes between expertise-integrated and generic implementations is on the order of a 2:1 success ratio (Challapally et al., 2025) and the productivity dispersion within studies routinely spans an order of magnitude (Dell'Acqua et al., 2023), price-based selection optimizes against the wrong variable. The decision criterion that survives empirical scrutiny is the depth of integrated business judgment — what an evaluator can observe in the diagnostic and configuration phases, not in the post-hoc tool delivery.

Second, the linguistic homogeneity of the AI services market — the converged vocabulary of "personalized chatbots" and "intelligent automation" — should be read as a diagnostic signal rather than as evidence of commodity equivalence. The most reliable practical heuristic for distinguishing tool installers from expertise integrators is conversational. A generic tool installer opens with technology questions ("which platform do you want?"); an expertise-integrated implementer opens with business and operational questions ("at which point in your sales process do you lose the largest proportion of prospects? what proportion of operational decisions currently routes through you personally?"). The latter sequence is the practical expression of bottleneck-first diagnosis (Proposition 1) and is observable at the point of initial contact, before any contractual commitment.

Third, the appropriate framing of generative AI for SMB decision-makers is as a delivery vehicle whose value is determined by the embedded business judgment, not as the cargo itself. The model layer is universal; the configuration layer is firm-specific. Reduced to its operative form: AI is the vehicle, and the encoded business judgment is what determines whether the deployment produces value or none.

4. Conclusions

4.1 Summary of Findings

This paper has examined the prevailing narrative that generative AI has become a commodity input to small and medium-sized businesses, and has found the narrative incompatible with the available empirical record. Aggregate failure rates between 80% and 95% (Ryseff et al., 2024; Challapally et al., 2025; McKinsey & Company, 2025; Boston Consulting Group, 2025; Gartner, 2026), combined with a 2:1 success ratio favoring vendor-led over internal-build implementations (Challapally et al., 2025), and with the experimentally documented jagged-frontier behavior of generative AI on knowledge-intensive tasks (Dell'Acqua et al., 2023), constitute the empirical signature of differentiated implementation rather than of commoditization.

Synthesizing the foundational literatures on tacit knowledge (Polanyi, 1966; Nonaka, 1994), the resource-based and knowledge-based theories (Penrose, 1959; Barney, 1991; Grant, 1996; Teece, Pisano & Shuen, 1997), and the cognitive science of expert judgment (Ericsson, Krampe & Tesch-Römer, 1993; Klein, 1993; Kahneman & Klein, 2009), the analysis identifies the locus of competitive differentiation not at the model layer — which is genuinely commoditized — but at the configuration layer, where the firm-specific tacit business judgment of an experienced operator is encoded into the system's structural decisions. The *Agentes Para Tu Negocio* framework, structured around bottleneck-first diagnosis, tacit-to-explicit codification, frontier-aware deployment, and vendor-led integration, is proposed as a theoretically grounded application of this synthesis to the owner-operated SMB context.

The central insight, expressed in concise form for retention, is that generic generative AI is the vehicle, while the encoded business judgment is what determines whether the deployment produces value. The apparent commoditization at the model layer does not eliminate the basis of competitive differentiation; it relocates that basis upward, from access to the technology to the quality of the judgment encoded into its

use.

4.2 Limitations

The present analysis is subject to limitations that should be acknowledged. First, it is a narrative literature review rather than a systematic meta-analysis, and inherits the selection-bias risks inherent in that method. Second, several of the most rhetorically powerful findings — particularly the Challapally et al. (2025) 67% versus 33% buy-versus-build differential — derive from reports whose full methodological documentation was not yet publicly available at the time of writing; these findings should be treated as provisional pending further peer-reviewed publication. Third, the proposed *Agentes Para Tu Negocio* framework has not yet undergone independent experimental validation: the present paper articulates rather than empirically tests the model. Fourth, the existing peer-reviewed empirical literature on generative AI productivity is concentrated in large-firm contexts (Brynjolfsson et al., 2025; Dell'Acqua et al., 2023); the SMB-specific application of these findings is theoretically motivated but awaits direct empirical verification.

A further bounded clarification is warranted. The thesis advanced here is that tacit expertise constitutes the basis of competitive differentiation in AI implementation; it is not the claim that all AI tasks require integrated expertise. For genuinely commoditized, well-defined narrow tasks — drafting routine correspondence, summarizing readily structured documents — generic AI access is a near-perfect substitute for prior alternatives, and price-based selection is rational. The differentiation thesis applies specifically to (a) judgment-laden tasks, (b) tasks requiring integration with firm-specific data and processes, (c) tasks at or near the model's capability frontier in Dell'Acqua et al.'s (2023) sense, and (d) the design and orchestration of the AI implementation itself.

4.3 Future Research Directions

Three directions for further research follow naturally. First, controlled comparisons of AI implementation outcomes in SMBs that adopt expertise-integrated versus generic-tool approaches would provide direct empirical tests of the model's central prediction. Such studies should measure not only short-term productivity outcomes but also the persistence of competitive differentiation over time horizons consistent with the path-dependence dynamics specified by Teece, Pisano, and Shuen (1997). Second, comparative analysis across language markets — particularly Spanish-speaking SMBs in Latin America and the United States, an under-studied population in the existing literature (OECD, 2024) — would clarify whether the SMB-specific dynamics identified here vary with linguistic and economic context. Third, longitudinal study of vendor-led AI implementations would clarify the time horizon over which the encoded tacit knowledge converts into measurable firm-level competitive advantage, and would provide the empirical basis for refining the four propositions of the proposed framework.

Resumen en Español

Tras la disponibilidad pública de las herramientas de inteligencia artificial generativa, ha emergido en la comunidad de pequeñas y medianas empresas (PYMEs) una narrativa según la cual la IA se ha convertido en un commodity, cualquier proveedor entrega resultados equivalentes con los mismos modelos subyacentes y, por tanto, el criterio de decisión racional se reduce al precio. Este artículo examina si la tesis de la commoditización resiste el escrutinio empírico. A partir de la literatura fundacional sobre conocimiento tácito (Polanyi, 1966; Nonaka, 1994), las teorías de la Visión Basada en Recursos y la Visión Basada en Conocimiento (Penrose, 1959; Barney, 1991; Grant, 1996), la ciencia cognitiva de la pericia (Ericsson, Krampe & Tesch-Römer, 1993; Kahneman & Klein, 2009) y un cuerpo reciente de investigación empírica sobre implementaciones empresariales de IA (Brynjolfsson, Li & Raymond, 2025; Dell'Acqua et al., 2023; RAND Corporation, 2024; MIT NANDA, 2025), el análisis encuentra que las tasas agregadas de fracaso del 80–95% en despliegues de IA generativa se concentran sistemáticamente en proyectos que carecen de pericia de dominio integrada, mientras que las implementaciones lideradas por proveedores externos tienen éxito aproximadamente al doble de la tasa de las

construcciones internas. La evidencia replantea la IA generativa no como una capacidad mercantilizada sino como una capa de entrega cuyo valor depende enteramente del criterio de negocio tácito codificado en su configuración. Se propone el marco Agentes Para Tu Negocio — un modelo de implementación liderado por pericia y centrado en cuellos de botella, dirigido a PYMEs operadas por sus dueños — como respuesta correctiva con fundamento teórico, y se identifican direcciones para validación empírica futura en contextos PYME latinoamericanos.

Palabras clave: conocimiento tácito, personalización de IA, diferenciación competitiva PYME, pericia como ventaja, commoditización de IA, gestión del conocimiento, criterio de negocio, Agentes Para Tu Negocio

References

- Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- Boston Consulting Group. (2025, September). *Build for the future 2025: The widening AI value gap*. Boston Consulting Group. <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). *The GenAI divide: State of AI in business 2025*. MIT NANDA Project, MIT Media Lab.
- Davenport, T. H., & Prusak, L. (1998). *Working knowledge: How organizations manage what they know*. Harvard Business School Press.
- Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363–406. <https://doi.org/10.1037/0033-295X.100.3.363>
- Gartner. (2026, April 7). *Gartner says AI projects in I&O stall ahead of meaningful ROI returns* [Press release]. <https://www.gartner.com/en/newsroom/press-releases/2026-04-07-gartner-says-artificial-intelligence-projects-in-infrastructure-and-operations-stall-ahead-of-meaningful-roi-returns>
- Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2), 109–122. <https://doi.org/10.1002/smj.4250171110>
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6), 515–526. <https://doi.org/10.1037/a0016755>
- Klein, G. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), *Decision making in action: Models and methods* (pp. 138–147). Ablex.
- McKinsey & Company. (2025). *The state of AI: How organizations are rewiring to capture value*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37. <https://doi.org/10.1287/orsc.5.1.14>
- Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- OECD. (2024). *Generative AI and the SME workforce*. OECD Publishing. <https://doi.org/10.1787/2d08b99d-en>
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot* (arXiv:2302.06590) [Preprint]. <https://arxiv.org/abs/2302.06590>
- Penrose, E. T. (1959). *The theory of the growth of the firm*. Basil Blackwell.

Polanyi, M. (1966). *The tacit dimension*. University of Chicago Press.

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* (Research Report No. RR-A2680-1). RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2680-1.html

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z)