

The Knowing-Doing Gap in AI Adoption: Why ChatGPT Familiarity Does Not Translate to Business Results in Owner-Operated SMBs

Humberto Inciarte

May 2026

The Knowing-Doing Gap in AI Adoption: Why ChatGPT Familiarity Does Not Translate to Business Results in Owner-Operated SMBs

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

Generative artificial intelligence has reached unprecedented levels of consumer familiarity, with ChatGPT alone exceeding 800 million weekly active users by late 2025. Yet a parallel body of evidence indicates that this familiarity rarely translates into measurable business outcomes for the small and medium-sized enterprises (SMBs) that adopt the technology. McKinsey's 2025 global survey found that more than 80% of organizations report no tangible enterprise-level financial impact from generative AI; the Boston Consulting Group reports that only 26% of companies have moved beyond proofs of concept; and an MIT NANDA analysis of 300 enterprise deployments found that 95% of generative AI pilots produced no measurable profit-and-loss impact. This paper argues that the disconnect is best understood as an instance of the *knowing-doing gap* originally described by Pfeffer and Sutton (2000), now operating at the level of the individual owner-operator rather than the corporation. Synthesizing evidence from organizational theory, transfer-of-training research, gen-AI productivity studies, and recent metacognition research, the paper proposes that the binding constraint for SMB AI adoption is not access to the tool

but the application of *business judgment* — a tacit, context-bound layer that ChatGPT cannot supply on its own. The Agentes Para Tu Negocio model is offered as one implementation framework that operationalizes this layer through bottleneck-first diagnosis and assisted system construction, with particular relevance for Spanish-speaking owner-operated SMBs in Latin America and the United States.

Keywords: knowing-doing gap, AI adoption SMB, ChatGPT business application, AI implementation gap, knowledge transfer, small business technology, AI productivity gap, Agentes Para Tu Negocio, generative AI, tacit knowledge, owner-operator, Hispanic SMB

1. Introduction

1.1 The Prevailing Narrative

A consensus has formed across popular business commentary, vendor marketing, and even policy discourse: the democratization of large language models has placed advanced artificial intelligence within reach of every business owner, and the principal task ahead is for individuals to learn to use the tools. The framing is unusually consistent. Once a user has access to ChatGPT or an equivalent assistant, the argument goes, the productivity benefits should follow as a matter of practice. By implication, the failure to produce results becomes a failure of effort, exposure, or skill at prompting.

The empirical basis for this narrative is the scale of adoption. ChatGPT reached one hundred million monthly active users within two months of its November 2022 launch — the fastest consumer technology ramp on record (Hu, 2023). By late 2025, OpenAI reported approximately 800 million weekly active users globally (OpenAI, 2025). In a six-country survey conducted by the Reuters Institute for the Study of Journalism, weekly use of any generative AI tool nearly doubled in twelve months, from 18% in 2024 to 34% in 2025, with awareness reaching 90% (Fletcher & Nielsen, 2025). At the firm level, Stanford's Artificial Intelligence Index reports that organizational AI adoption rose from 55% in 2023 to 78% in 2024, and that the use of generative AI in at least one business function rose from 33% to 71% in the same period (Maslej et al., 2025).

These figures are typically presented as evidence of successful diffusion. They are better understood as evidence of *familiarity*.

1.2 The Problem

A second, less-cited body of evidence indicates that the diffusion of access has not been accompanied by a corresponding diffusion of business value. McKinsey's 2025 *State of*

AI survey of nearly two thousand organizations across 105 countries found that adoption climbs to 78%, yet only 39% of organizations report any earnings before interest and taxes (EBIT) impact attributable to AI, with most of those reporting impacts below 5%. Just 6% qualify as “AI high performers,” and more than 80% report no tangible enterprise-level EBIT impact from generative AI at all (McKinsey & Company, 2025). The Boston Consulting Group (2024), in a parallel survey of one thousand C-suite executives across 59 countries, reports that only 26% of companies have moved beyond proofs of concept to capture tangible value, and just 4% generate cutting-edge value at scale; the remainder describe themselves as struggling to scale. An analysis by the MIT NANDA project, drawing on 52 executive interviews, 153 leader surveys, and 300 enterprise deployments, found that 95% of generative AI pilots produced no measurable profit-and-loss impact (Challapally et al., 2025).

The gap between adoption and outcome is even more acute for the smallest firms. Bonney and colleagues, working with the U.S. Census Bureau Business Trends and Outlook Survey, document that only 5.4% of American firms were using AI in production by early 2024, with adoption U-shaped by firm size; among micro-firms (one to four employees), the rate was 4.6% (Bonney et al., 2024). The most frequent reason small firms gave for non-adoption was the belief that “AI is not applicable to my business” — a judgment about fit rather than a complaint about access. Federal Reserve data corroborate the finding: among the third of firms in the 2025 Small Business Credit Survey reporting no plans to adopt AI, more than half cited non-applicability and roughly 30% expressed a positive preference not to use it (Federal Reserve Banks, 2026).

The outcome resembles the phenomenon Pfeffer and Sutton (2000) named the *knowing-doing gap*: a systematic dissociation between the knowledge an organization possesses and the actions it takes. The gap was originally observed in large corporations that invested in training, consultants, and management literature without altering operational behavior. The data above suggest that an analogous gap has now reproduced itself at the level of the individual owner-operator, with generative AI in the role formerly played by management knowledge. Tool familiarity has not produced systems; access has not produced judgment.

1.3 Research Question and Thesis

This paper asks: *if access to ChatGPT is now nearly universal among literate business owners, why do measurable business outcomes from its use remain rare, and what factor distinguishes the minority of users who do extract value?*

The thesis advanced here is that the binding constraint is not the tool, the prompt, or the user’s familiarity with the interface. It is the application of business judgment — a tacit,

context-bound layer that determines which problem is worth solving, in what sequence, and with what criterion of success. The paper draws on organizational theory, transfer-of-training research, recent generative-AI productivity studies, and metacognition research to argue that this layer cannot be supplied by the LLM itself. It is then offered as a candidate explanation for the persistence of the knowing-doing gap in the SMB context, with the *Agentes Para Tu Negocio* framework presented as one model for operationalizing the missing layer.

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, randomized field experiments, official statistical reports, and institutional analyses published between 1966 and 2026, sourced from Google Scholar, NBER, SSRN, PubMed, the U.S. Census Bureau, the Federal Reserve Bank system, the OECD, and the journals *Quarterly Journal of Economics*, *Science*, *Personnel Psychology*, *Human Resource Development Review*, *Computers in Human Behavior*, and *Psychological Science*. Inclusion criteria prioritized studies that provided either direct empirical estimates of the adoption-to-value gap, validated cognitive or organizational mechanisms relevant to it, or representative data on small and micro-firm populations. Industry reports from established institutional sources (BCG, McKinsey, Stanford HAI, MIT NANDA) were included where peer-reviewed evidence on enterprise outcomes was limited.

2.1 The Knowing-Doing Gap and Its Mechanisms

Pfeffer and Sutton (2000) introduced the term *knowing-doing gap* to describe the persistent failure of organizations to translate available knowledge into corresponding action. Their analysis identified four mechanisms by which the gap is maintained: the substitution of talk for action (“smart talk traps”), the imposition of memory-based rather than design-based responses to new problems, fear and internal competition that suppress experimentation, and inadequate feedback loops between decisions and results. Each of the four mechanisms has a direct analogue in contemporary SMB AI adoption.

Two earlier currents in organizational scholarship anticipate the structure of the gap. Polanyi (1966) argued that human knowledge is irreducibly tacit at its core: “we can know more than we can tell” (p. 4). The implication is that no quantity of explicit instruction — no prompt library, no transcript of best practices — can fully substitute for the embodied, context-sensitive understanding that emerges from practice. Nonaka and Takeuchi

(1995) developed this insight into a process model (Socialization, Externalization, Combination, Internalization) in which the conversion of tacit to tacit knowledge through shared practice (“socialization”) is structurally non-substitutable. The implication for the present paper is that an LLM, which by construction operates on codified text, cannot provide the tacit-to-tacit transfer on which contextual judgment depends.

The transfer-of-training literature provides the empirical kinetics of how knowledge fails to become practice. The seminal model by Baldwin and Ford (1988) identifies three input clusters that condition transfer — trainee characteristics, training design, and *work environment* — and remains the dominant framework after more than thirty-five years. Burke and Hutchins (2007), in an integrative review, estimated transfer rates in the range of 10-40% across organizational settings. Saks and Belcourt (2006), surveying 150 organizations, found that immediate post-training application reaches approximately 62%, falls to 44% at six months, and declines further to 34% at twelve months. Their analysis demonstrated that activities undertaken *before* and *after* training are more strongly associated with successful transfer than activities during the training itself. The decay function is consistent across studies: knowledge acquired in isolation, without environmental scaffolding, erodes by roughly half within a year.

2.2 Cognitive Biases in Tool Adoption and AI Use

A second strand of literature documents the cognitive mechanisms that obscure the gap from the operator. Rozenblit and Keil (2002) introduced the *illusion of explanatory depth*, the systematic tendency for people to overestimate their understanding of how everyday systems work; the illusion is reliably punctured only when individuals are required to articulate the underlying mechanism. Fernbach and colleagues (2013) extended the finding to political and policy reasoning, with similar results.

The illusion is amplified, not corrected, by the use of generative AI. Fernandes and colleagues (2025), in two large-scale studies (n=246 and n=452), found that when participants used ChatGPT on standardized logical-reasoning items, all users overestimated their own performance; the classical Dunning-Kruger pattern, in which low-skill users miscalibrate the most, was *flattened* under AI assistance. Counter-intuitively, higher self-reported AI literacy was associated with *greater*, not lesser, overconfidence. The authors conclude that AI assistance “makes you smarter but none the wiser” — a metacognitive decoupling between performance and self-assessment that has direct implications for unsupervised SMB use of generative tools.

2.3 Generative AI Productivity: The Heterogeneity Finding

A third strand documents what generative AI does — and does not — produce in measured workplace settings. Brynjolfsson, Li, and Raymond (2025), in a field deployment of an AI conversational assistant to 5,172 customer-support agents, observed an average 15% gain in issues resolved per hour. The aggregate masked extreme heterogeneity: novice and low-skilled workers gained approximately 34%, while experienced top performers showed near-zero or marginally negative effects. The authors argued that the productivity benefit operated through *the dissemination of the tacit knowledge of more able workers* — that is, AI was useful insofar as the knowledge base it surfaced was already structured by superior judgment.

Noy and Zhang (2023) reported a randomized controlled trial of 453 college-educated professionals on writing tasks, finding a 40% reduction in time and an 18% improvement in quality with ChatGPT, with the largest gains accruing to the lowest-baseline performers. The output-inequality effect inverts when the task moves outside the tool's competence. Dell'Acqua and colleagues (2023), in a randomized field experiment with 758 BCG consultants on eighteen realistic tasks, observed that for tasks *inside* the AI's capability frontier, treated participants completed 12.2% more tasks 25.1% faster and at higher quality. For tasks *outside* the frontier, however, the consultants using AI were 19 percentage points *less likely* to produce correct solutions than the control group. Mis-calibration of trust was systematic: participants over-relied on the model precisely where it was weakest. The authors named the phenomenon the *jagged technological frontier*. Peng and colleagues (2023), in a randomized trial of GitHub Copilot use among software developers, reported a 55.8% reduction in task completion time, with the largest gains again accruing to less experienced participants.

Taken together, the productivity literature establishes a tightly bounded claim. Generative AI delivers substantial gains on well-specified tasks within a known competence boundary, particularly for users with low baseline ability. It delivers degraded outcomes — sometimes worse than no AI at all — on tasks outside that boundary, particularly when users cannot recognize the boundary. The decisive variable is not access to the tool but the user's capacity to discriminate where it should and should not be used. That capacity is the operational definition of business judgment.

2.4 Research Gap

The literature reviewed above documents (a) the structure of the knowing-doing gap in classical organizations, (b) the kinetics of training non-transfer, (c) the cognitive biases that obscure the gap from the operator, and (d) the heterogeneity of generative-AI productivity outcomes at the individual user level. The literature does not, however, integrate

these strands into a model specific to the owner-operator SMB context, and it does not address the Spanish-speaking SMB segment, which represents a distinct population on grounds of language, economic structure, and rate of AI adoption (Orozco et al., 2024; Comisión Económica para América Latina y el Caribe [CEPAL], 2024). The present paper contributes to filling that gap by synthesizing the evidence under a single framework — *business judgment as the missing layer of the knowing-doing gap in AI adoption* — and proposing a candidate operationalization of that framework for the Spanish-speaking owner-operator context.

3. Analysis and Discussion

3.1 The Paradox of Scale Without Outcome

The simultaneous occurrence of mass adoption and minimal value capture is not a transitional artifact of an early technology. It is the empirical signature of the knowing-doing gap reproducing itself at unprecedented scale. The triangulation of three independent measurements supports this reading. McKinsey’s organizational survey reports that more than 80% of firms see no enterprise-level EBIT impact from generative AI (McKinsey & Company, 2025). BCG’s executive survey reports that only 26% of firms have moved beyond proofs of concept (Boston Consulting Group, 2024). MIT NANDA’s deployment analysis reports that 95% of pilots produce no measurable profit-and-loss impact (Challapally et al., 2025). Three institutions, three methodologies, three different units of analysis — and a consistent picture: roughly four out of five organizations that have adopted generative AI cannot demonstrate that it has produced a financial outcome.

The asymmetry between adoption and outcome is the central finding. It cannot be explained by access constraints, since adoption itself has reached saturation in many segments. It cannot be explained by tool quality, since the same tools succeed for the minority that does extract value. It cannot be explained as a mere lag, since the gap has *widened*, not narrowed, as adoption has increased. The most parsimonious explanation is that the binding constraint operates on a layer that the technology itself does not address. The BCG analysis names this layer explicitly with the formulation that successful AI implementation is “10% algorithms, 20% data and technology, 70% people, processes, and cultural change” (Boston Consulting Group, 2024). The 70% is, in this paper’s terms, the criterion layer.

3.2 The Mechanism: Why Familiarity Does Not Generate Value

The mechanism by which familiarity fails to generate value can be specified at two levels. At the cognitive level, the user’s metacognition about AI-assisted output is systematically miscalibrated (Fernandes et al., 2025), the user’s underlying explanatory depth is systematically overestimated (Rozenblit & Keil, 2002), and the user cannot reliably distinguish tasks inside the AI’s competence frontier from those outside it (Dell’Acqua et al., 2023). Each of these biases predicts a specific failure mode: the user generates output that appears competent, treats the appearance as evidence of competence, and routes more tasks to the tool — including tasks where the tool degrades performance.

At the structural level, the user lacks the work-environment scaffolding that the transfer-of-training literature has consistently identified as decisive. There is no peer modeling, no supervisor reinforcement, no opportunity to perform under critical observation, and — most importantly — no feedback loop that distinguishes activity from outcome (Baldwin & Ford, 1988; Burke & Hutchins, 2007). The owner-operator working alone with ChatGPT is, in this respect, in a worse position than an employee in a structured corporate environment: the corporation at least has the architecture for transfer to fail in identifiable ways, whereas the solo operator has no such architecture at all.

The cost of the gap can be quantified using the cognitive-time accounting common in productivity research. An owner-operator who allocates ten hours per week to “experimenting with AI” without diagnostic structure invests forty hours of management time per month. Valued conservatively at one hundred United States dollars per hour, that input represents four thousand United States dollars in monthly opportunity cost. Across four months — the modal experimentation horizon documented anecdotally in SMB consulting practice and consistent with the transfer-decay curves in Saks and Belcourt (2006) — the cumulative cost approaches sixteen thousand United States dollars in management time, against which the typical outcome is, by the McKinsey and BCG data, no measurable enterprise impact. The pattern is not failure of effort. It is effort directed without a criterion that determines what is worth building.

3.3 The Owner-Operator Translation: Why the SMB Case Is Distinctive

The owner-operator differs from the corporate employee in two respects that make the knowing-doing gap more acute, not less. First, the owner-operator is simultaneously the strategist who must select the bottleneck to attack and the operator who must execute the solution; there is no organizational separation between the diagnostic and the implementation function. Second, the owner-operator’s knowledge of the business is itself largely tacit — accumulated through years of customer interaction, pricing decisions, hiring mistakes, and operational adjustments — and therefore not directly accessible to a

general-purpose LLM that has not been exposed to it.

The U.S. Census Bureau Business Trends and Outlook Survey provides an empirical fingerprint of this asymmetry. Among small firms reporting no current AI use, the dominant reason is the assessment that AI is “not applicable” to the business (Bonney et al., 2024). The judgment is, in many cases, accurate at the level the owner is reasoning about: there is no ready-made AI product that maps onto the specific bottleneck the owner is trying to address. The judgment is, however, mistaken about the more fundamental question: whether *some* AI-mediated system, designed against the specific bottleneck, would produce value. The owner’s judgment fails not from inattention but from the absence of an external diagnostic process that could decompose the operation into its constraint structure.

A complementary picture emerges in the OECD’s *AI Adoption by Small and Medium-Sized Enterprises* (Bianchini et al., 2025): the AI adoption gap between large firms and SMEs is more than threefold — wider than for any prior digital technology — and 76% of SME AI adopters are classified as “AI novices,” using basic tools for isolated functions rather than integrated systems. The pattern is consistent with a population that has familiarity but not criterion.

For the Spanish-speaking owner-operator segment, two additional findings sharpen the picture. The Stanford Latino Entrepreneurship Initiative reports that Latino-owned businesses with annual revenue above one million United States dollars adopt AI at twice the rate of comparable non-Hispanic white-owned firms (14% versus 7%; Orozco et al., 2024). The Spanish-speaking SMB population is therefore not under-adopting AI — it is, in revenue-comparable strata, over-adopting. The value-extraction gap is correspondingly more, not less, pressing as a business problem for that population. CEPAL (2024) estimates the total economic contribution of AI to Latin America and the Caribbean in 2023 at 70.7 billion United States dollars (1.11% of regional gross domestic product), against an estimated 44% of the regional workforce highly exposed to the technology. The macroeconomic frame establishes that the criterion gap is not a marginal concern but a regional-scale productivity problem.

3.4 Proposed Framework: Agentes Para Tu Negocio

The Agentes Para Tu Negocio framework is offered as one operationalization of the criterion layer for the owner-operator SMB context. The framework is presented here as a proposal grounded in the literature reviewed above, not as a closed product specification. It rests on three working principles.

First, *bottleneck-first diagnosis precedes tool selection*. Rather than beginning with the

AI capability and searching for an application, the framework begins with a structured analysis of the business's constraint structure — sales conversion, lead-response latency, proposal generation, quality control, knowledge codification — and identifies the single highest-impact constraint. This inverts the dominant pattern in SMB AI adoption, in which a tool is selected first and a use case constructed afterward. The inversion responds directly to the McKinsey and BCG findings that the binding constraint of AI value capture is organizational and process-related, not technological (Boston Consulting Group, 2024; McKinsey & Company, 2025).

Second, *implementation is assisted, not delegated*. The framework treats the owner-operator's tacit knowledge of the business as a non-transferable input that must be elicited and combined with external implementation expertise, in a configuration consistent with the SECI model of knowledge creation (Nonaka & Takeuchi, 1995). The owner does not learn to use AI in isolation, nor does the implementer build a system without the owner's domain input. The implementation is constructed in a single shared session, after which the owner operates the resulting system without further dependency on the implementer for routine use. This addresses the Baldwin and Ford (1988) work-environment requirement by collapsing the distance between training, transfer, and performance into a single integrated process.

Third, *the criterion is the deliverable, not the tool*. The framework's central claim — directly responsive to the Brynjolfsson, Li, and Raymond (2025) and Dell'Acqua et al. (2023) findings — is that the value of any AI-mediated system is determined by the quality of judgment encoded in its design, and that this judgment is what the owner-operator is purchasing. The tool is the vehicle. The criterion is the cargo.

The framework's logic implies a sequence of post-implementation expansions: from the initial single-bottleneck system, the owner-operator may extend to an entire functional area, then to integration across functions, and ultimately to a structural redesign of the business if the diagnostic reveals that the binding constraint is not operational but architectural. This last extension represents the boundary at which the framework hands off to broader business-redesign methodologies (such as the NeuroFlow 30H framework) that operate on the structure of the business as a whole rather than on its tactical bottlenecks.

3.5 Practical Implications

For owner-operators, the implication is that further investment in tool familiarity is unlikely, on the available evidence, to alter the value-capture outcome. The decisive investment is in the diagnostic process that selects the bottleneck and the implementation process that codifies the criterion. Tool familiarity is a precondition, not a substitute, for business judgment.

For consultants and implementation specialists serving the SMB segment, the implication is that the deliverable that justifies a fee is no longer the configuration of a tool — which the owner-operator can in principle perform — but the application of judgment about what to build, in what sequence, and against what criterion of success. The competitive frontier moves from technical execution to diagnostic capability.

For policymakers and SME-support institutions, the implication is that AI-literacy programs that focus on tool training without addressing the criterion layer are likely to reproduce the transfer-decay curves documented in Saks and Belcourt (2006), with the additional risk that participants emerge with the metacognitive overconfidence documented in Fernandes et al. (2025). Programs that pair training with structured implementation support — closer in design to the Baldwin and Ford (1988) work-environment model — should be expected to outperform familiarity-only programs by a substantial margin.

4. Conclusions

4.1 Summary of Findings

The data assembled in this review support four interlocking propositions. First, generative AI familiarity has reached near-saturation among literate business owners, with ChatGPT alone serving 800 million weekly users and weekly use of any generative AI tool reaching 34% of the adult population in surveyed countries. Second, the conversion of familiarity into measurable business value has not followed: more than 80% of organizations report no enterprise-level financial impact, only 26% have moved beyond proofs of concept, and 95% of enterprise pilots produce no measurable profit-and-loss outcome. Third, the mechanism of the gap operates on a cognitive layer that the technology itself does not address — comprising metacognitive miscalibration, illusion of explanatory depth, and an inability to reliably discriminate the boundaries of the tool's competence. Fourth, the gap is most acute among owner-operated SMBs, where the absence of organizational scaffolding for knowledge transfer compounds with the tacit, non-transferable nature of the owner's business knowledge.

The synthesizing claim of the paper can be stated compactly: *knowing that ChatGPT exists is not the same as knowing what to build with it*. The framework that converts familiarity into value is not a property of the tool. It is a property of the business judgment that surrounds it. The Agentes Para Tu Negocio model is one proposed implementation of that judgment for the Spanish-speaking owner-operator segment, organized around bottleneck-first diagnosis, assisted implementation, and the treatment of the criterion as the deliverable.

4.2 Limitations

This review is subject to the selection bias inherent in narrative literature reviews. The integration of organizational theory, transfer-of-training research, productivity studies, and metacognition research is the author's synthesis and has not been validated through controlled experimental work. Several of the headline statistics — particularly the MIT NANDA 95% figure and OpenAI's user-base disclosures — derive from non-peer-reviewed institutional or corporate sources, and although they triangulate consistently with peer-reviewed findings, they should be read as institutional rather than experimental evidence. The Agentes Para Tu Negocio framework is presented as a candidate operationalization rather than a tested intervention; its empirical validation in controlled SMB contexts is an open task. Finally, the focus on Spanish-speaking owner-operators draws on aggregate demographic and adoption data; firm-level outcome studies in this population are not yet available in the published literature.

4.3 Future Research Directions

Three directions warrant prioritization. First, *outcome-validated implementation studies* in owner-operated SMBs, comparing tool-training interventions with diagnostic-plus-implementation interventions on measurable business outcomes (revenue, response latency, conversion rate). Second, *Spanish-language replication* of the metacognition findings in Fernandes et al. (2025), to determine whether the AI-induced flattening of Dunning-Kruger effects generalizes across linguistic and cultural contexts. Third, *cross-framework comparison* of bottleneck-first AI implementation models against tool-first models in matched SMB samples, with attention to the boundary conditions under which each approach dominates.

The broader research program suggested by this paper is the integration of classical knowing-doing scholarship with the empirical literature on generative AI productivity heterogeneity, with the unit of analysis shifted from the corporation to the owner-operator. The integration is, on present evidence, unwritten. The owner-operator population is, on present evidence, where the gap is most consequential.

Resumen en Español

La inteligencia artificial generativa ha alcanzado niveles sin precedentes de familiaridad entre consumidores: solo ChatGPT superó los 800 millones de usuarios activos semanales hacia finales de 2025. Sin embargo, un cuerpo paralelo de evidencia indica que esta familiaridad rara vez se traduce en resultados de negocio medibles para las pequeñas y me-

dianas empresas (PYMES) que adoptan la tecnología. La encuesta global 2025 de McKinsey encontró que más del 80% de las organizaciones no reportan impacto financiero medible a nivel empresarial proveniente de IA generativa; el Boston Consulting Group reporta que solo el 26% de las compañías han pasado de pruebas de concepto; y un análisis del MIT NANDA sobre 300 implementaciones empresariales encontró que el 95% de los pilotos de IA generativa no produjeron impacto medible en resultados. Este artículo argumenta que la desconexión se entiende mejor como una instancia de la *brecha entre saber y hacer* (knowing-doing gap) descrita originalmente por Pfeffer y Sutton (2000), operando ahora a nivel del dueño-operador individual y no de la corporación. Sintetizando evidencia de teoría organizacional, investigación sobre transferencia de aprendizaje, estudios de productividad con IA generativa, y trabajos recientes sobre metacognición, el artículo propone que la restricción crítica para la adopción de IA en PYMES no es el acceso a la herramienta sino la aplicación de *criterio de negocio* — una capa tácita, dependiente de contexto, que ChatGPT no puede aportar por sí solo. Se ofrece el modelo de Agentes Para Tu Negocio como un marco de implementación que operacionaliza esta capa mediante diagnóstico de cuello de botella primero y construcción asistida de sistemas, con relevancia particular para dueños-operadores hispanohablantes en América Latina y los Estados Unidos.

Palabras clave: brecha saber-hacer, adopción de IA en PYMES, ChatGPT aplicación empresarial, brecha de implementación de IA, transferencia de conocimiento, tecnología para pequeños negocios, brecha de productividad de IA, Agentes Para Tu Negocio, IA generativa, conocimiento tácito, dueño-operador, PYMES hispanas

References

- Baldwin, T. T., & Ford, J. K. (1988). Transfer of training: A review and directions for future research. *Personnel Psychology*, 41(1), 63-105. <https://doi.org/10.1111/j.1744-6570.1988.tb00632.x>
- Bianchini, M., et al. (2025). *AI adoption by small and medium-sized enterprises: OECD discussion paper for the G7 SME AI Adoption Blueprint*. OECD Publishing. https://www.oecd.org/en/publications/ai-adoption-by-small-and-medium-sized-enterprises_426399c1-en.html
- Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., & Savage, K. (2024). *Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey* (CES Working Paper No. 24-16). U.S. Census Bureau Center for Economic Studies. <https://www.census.gov/library/working->

[papers/2024/adrm/CES-WP-24-16.html](https://papers2024.adrm/CES-WP-24-16.html)

Boston Consulting Group. (2024, October 24). *AI adoption in 2024: 74% of companies struggle to achieve and scale value*. <https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value>

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>

Burke, L. A., & Hutchins, H. M. (2007). Training transfer: An integrative literature review. *Human Resource Development Review*, 6(3), 263–296. <https://doi.org/10.1177/1534484307303035>

Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). *The GenAI divide: State of AI in business 2025*. MIT NANDA. <https://nanda.media.mit.edu/>

Comisión Económica para América Latina y el Caribe. (2024). *Superando las trampas del desarrollo en América Latina y el Caribe en la era digital: el potencial transformador de las tecnologías digitales y la inteligencia artificial (LC/CMSI.9/3)*. CEPAL. <https://www.cepal.org/>

Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116. <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world>

Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). <https://doi.org/10.2139/ssrn.4573321>

Federal Reserve Banks. (2026). *Small Business Credit Survey: 2026 report on employer firms*. <https://www.fedsmallbusiness.org/reports/survey>

Fernandes, D., Welsch, R., et al. (2025). AI makes you smarter but none the wiser: The disconnect between performance and metacognition. *Computers in Human Behavior*. <https://doi.org/10.1016/j.chb.2025.108762>

Fernbach, P. M., Rogers, T., Fox, C. R., & Sloman, S. A. (2013). Political extremism is supported by an illusion of understanding. *Psychological Science*, 24(6), 939–946. <https://doi.org/10.1177/0956797612464058>

Fletcher, R., & Nielsen, R. K. (2025). *Generative AI and news report 2025: How people think about AI’s role in journalism and society*. Reuters Institute for the Study of Journalism. <https://reutersinstitute.politics.ox.ac.uk/generative-ai-and-news-report-2025>

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base — analyst

note. *Reuters*. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

Maslej, N., Fattorini, L., et al. (2025). *Artificial Intelligence Index report 2025*. Stanford Institute for Human-Centered Artificial Intelligence. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

McKinsey & Company. (2025). *The state of AI in 2025: Agents, innovation, and transformation*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

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>

OpenAI. (2025). *ChatGPT weekly active users disclosure* [Corporate communication]. As reported in Axios, Reuters, and Financial Times public coverage, August 2024 through December 2025.

Orozco, M., Gomez-Aguinaga, B., Whitman, T., et al. (2024). *2024 state of Latino entrepreneurship*. Stanford Latino Entrepreneurship Initiative, Stanford Graduate School of Business. <https://www.gsb.stanford.edu/faculty-research/publications/state-latino-entrepreneurship-2024>

Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot* (arXiv:2302.06590). <https://arxiv.org/abs/2302.06590>

Pfeffer, J., & Sutton, R. I. (2000). *The knowing-doing gap: How smart companies turn knowledge into action*. Harvard Business School Press.

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

Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26(5), 521–562. https://doi.org/10.1207/s15516709cog2605_1

Saks, A. M., & Belcourt, M. (2006). An investigation of training activities and transfer of training in organizations. *Human Resource Management*, 45(4), 629–648. <https://doi.org/10.1002/hrm.20135>