

Guided Implementation vs. Self-Directed Adoption: Comparative Outcomes of Expert-Mediated AI Integration in Hispanic SMBs

Humberto Inciarte

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

ORCID: [0009-0006-0379-7328](https://orcid.org/0009-0006-0379-7328)

hola@humbertoinciarte.com

May 2026

Abstract

The dominant narrative around generative AI adoption in small and medium-sized businesses treats the technology itself as the primary determinant of outcome. Under this framing, the rational owner-operator concludes that hiring an external implementer is an avoidable expense — a cost made visible by an invoice and weighed against tools that appear nearly free. This paper examines whether that conclusion holds up under empirical scrutiny. Synthesizing evidence from six independent research streams — executive coaching meta-analyses, MOOC attrition data, SaaS onboarding benchmarks, the Implementation Science literature on facilitation, recent enterprise GenAI deployment research, and educational psychology on scaffolding and deliberate practice — the analysis reveals a consistent pattern: expert-mediated implementation outperforms self-directed adoption by ratios approaching two to one across measurable outcomes. The 2025 MIT NANDA *State of AI in Business* report finds that vendor-partnered AI deployments succeed roughly 67% of the time versus 33% for internal self-directed builds. Comparable ratios appear in SCORE small-business mentoring data, SaaS activation benchmarks, and decades of facilitation research in health services. The paper introduces the *Agentes Para Tu Negocio* model — a “done-WITH-you” implementation approach situated within the Implementation Facilitation construct established by Stetler, Kirchner, and colleagues — as a structurally appropriate response for Hispanic-owned SMBs, a population whose AI adoption rate doubled between 2024 and 2025 yet remains underserved by the existing dichotomy of pure DIY or opaque done-FOR-you outsourcing. The proposed model treats business judgment, not the tool, as the active ingredient.

1. Introduction

1.1 The Prevailing Narrative

The contemporary discourse on small business AI adoption rests on a quietly powerful assumption: that the value of a generative AI tool resides primarily in the tool itself. By this logic, what separates a successful adopter from an unsuccessful one is access, prompt quality, or hours invested. The owner-operator who is willing to read documentation, watch tutorials, and experiment with ChatGPT, Claude, or any of the dozens of competing platforms ought, in principle, to arrive at the same outcomes that an expensive consultant would deliver — only later, and at lower cash cost.

This framing is reinforced by three structural features of the current market. First, generative AI tools carry near-zero marginal cost at the point of access; ChatGPT Plus is approximately twenty

dollars per month, while a structural consulting engagement may cost thousands. Second, the consultant fee is visible on an invoice, while the cost of a failed self-implementation — months of unstructured experimentation, abandoned automations, decisions deferred — accrues invisibly across the calendar. Third, the most-shared narratives about AI in popular media celebrate solo operators who built remarkable systems alone, creating a survivorship-biased reference class.

The result is a near-universal default among owner-operators: do it yourself first, and only consider help if and when self-implementation visibly fails. This default is treated, by the actors involved, as the rational and prudent path. It is the position the authors of this paper set out to test.

1.2 The Problem

The empirical record suggests that the DIY default produces poor outcomes at scale. The 2025 MIT NANDA project's *State of AI in Business* report, drawing on 150 leader interviews, a survey of 350 employees, and analysis of 300 enterprise deployments, found that approximately 95% of generative AI pilots fail to deliver measurable profit-and-loss impact (Challapally et al., 2025). The report's most consequential finding for present purposes is the disaggregation: internally built, self-directed systems succeed roughly 33% of the time, while vendor-partnered or externally mediated deployments succeed approximately 67% of the time — a two-to-one differential. The report attributes the gap not to tool capability, which is essentially equivalent across both groups, but to what its authors term a “learning gap”: the absence of an external party who adapts the tool to the organization's actual workflow.

This pattern repeats outside the AI context. SCORE, the largest small-business mentoring organization in the United States, reports that small businesses receiving five or more mentoring interactions show substantially higher growth rates than those receiving none, with mentored businesses surviving past five years at roughly twice the rate of unmentored counterparts (SCORE, 2018). Stanford's *State of Latino Entrepreneurship* (Orozco et al., 2025) documents that Latino-owned businesses doubled their AI adoption rate between 2024 and 2025, yet face persistent structural asymmetries — only 21% of Latino entrepreneurs receive their full requested funding compared to 40% of white entrepreneurs — that compound the cost of failed self-implementation.

The problem, in short, is that the DIY default is producing a generation of owner-operators who experiment with AI, fail to capture business value, and conclude that “AI does not work for businesses like mine” — when the operative variable was never the tool.

1.3 Research Question and Thesis

This paper examines whether the DIY framing holds up under empirical scrutiny, and proposes an alternative model — *Agentes Para Tu Negocio* — that addresses the structural root causes identified in the literature. The central research question is: *across comparable populations and tasks, what is the measurable outcome differential between self-directed technology adoption and expert-*

mediated technology adoption, and what mechanism explains the gap?

The thesis advanced is that the differential is real, replicated across at least six research domains, and structurally explained by the absence of expert judgment — what the authors will refer to throughout as *criterio* — during the critical application phase of implementation. The paper proceeds as follows: Section 2 reviews the relevant literature across coaching, self-paced learning, SaaS onboarding, implementation science, enterprise AI adoption, and educational psychology. Section 3 analyzes the consistent two-to-one pattern that emerges, examines the cognitive biases that make DIY appear rational despite this evidence, and introduces the proposed framework. Section 4 concludes with implications, limitations, and directions for future research.

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, organizational case reports, and institutional data published between 1984 and 2026, sourced from PubMed, PsycINFO, Google Scholar, SSRN, the Stanford Latino Entrepreneurship Initiative archive, and the Implementation Science journal database. Inclusion criteria prioritized studies measuring the comparative performance of guided versus unguided participants on outcomes related to technology adoption, skill acquisition, or organizational implementation. Grey literature from established industry sources — McKinsey & Company, Boston Consulting Group, the MIT NANDA project, and the International Coaching Federation — was included where peer-reviewed evidence was limited but methodology was disclosed.

2.1 The Coaching and Guided-Learning Effect

The most rigorous body of evidence on the comparative outcomes of expert-mediated versus self-directed development comes from the coaching literature. The Theeboom, Beersma, and van Vianen (2014) meta-analysis aggregated effects across coaching studies in organizational contexts and reported significant positive effects on five outcome categories: performance and skills, well-being, coping, work attitudes, and goal-directed self-regulation. Hedges' g values ranged from 0.43 for coping to 0.74 for goal-directed self-regulation, with an aggregate effect size of approximately 0.66. Within-subjects designs produced an even larger g of 1.15.

Jones, Woods, and Guillaume (2016), publishing in the *Journal of Occupational and Organizational Psychology*, found across seventeen studies that workplace coaching positively affected organizational outcomes overall ($\delta = 0.36$), with skill-based outcomes at 0.28, affective outcomes at 0.51, and individual-level results reaching $\delta = 1.24$. Notably, the analysis found no significant moderating effect of delivery format — face-to-face versus e-coaching produced statistically equivalent results — suggesting that the active ingredient is the structured presence of an expert rather than the modality through which expertise is delivered.

The most-cited industry estimate of coaching return on investment, the Manchester study (Mc-

Govern et al., 2001), reported approximately 5.7 times return on initial investment based on perceptual estimates from forty-three of one hundred coached executives. This figure must be cited with care: as Ely and colleagues (2010) observed in their critical review, the methodology relied on retrospective self-report and was not a controlled trial. Nonetheless, the directional finding is consistent with the meta-analytic evidence above and with the longitudinal industry data from the International Coaching Federation, whose 2023 and 2025 *Global Coaching Studies* documented sustained growth of the practitioner population to nearly 123,000 worldwide and total industry revenue exceeding five billion dollars (International Coaching Federation, 2023, 2025).

2.2 The Self-Directed Cliff

A parallel literature documents what happens when expert mediation is removed. The earliest large-sample work in massive open online courses (MOOCs) by Jordan (2014, 2015) reported median completion rates of 12.6% across 221 courses, with a substantial subset falling below 5%. Reich and Ruipérez-Valiente (2019), publishing in *Science*, analyzed five years of MIT and Harvard edX courses and found that completion among all enrolled participants fell to 3.13% in the 2017–2018 academic year, down from 6% in 2014–2015. Among “verified” learners — those who paid for a certificate, the highest-intent self-directed cohort — completion was 46% in the most recent year measured, down from 56%. Most strikingly, 52% of registrants never started the course at all.

The pattern extends into the consumer software domain. Userpilot’s multi-year SaaS benchmark, drawn from sixty-two B2B platforms, found an average user activation rate of 37.5%, meaning roughly two-thirds of users who signed up never reached the product’s stated value moment (Userpilot, 2023). Industry retention data popularized through Localytics’ mobile analytics — the so-called 77/90/95 curve — indicates that the average application loses 77% of daily active users within three days of install, 90% within thirty days, and 95% within ninety days. While these figures are vendor-reported and lack the rigor of peer review, their consistency across reports and product categories establishes a reliable directional finding: pure self-service adoption produces a steep engagement decay curve.

This curve is what the authors term the *cliff of disengagement*. It does not describe people who lack motivation; it describes people whose motivation is real but whose first encounter with friction in an unfamiliar system has no scaffold to absorb it. The user is not lazy. The user is alone.

2.3 The Theoretical Foundation: Scaffolding, Cognitive Apprenticeship, and Deliberate Practice

The empirical findings above sit on a substantial theoretical foundation. Vygotsky’s (1978) construct of the Zone of Proximal Development posits that meaningful learning occurs in the space between what a learner can accomplish alone and what they can accomplish with expert guidance — a space that, by definition, is invisible to the unguided learner. Collins, Brown, and

Newman (1989), in their foundational work on cognitive apprenticeship, operationalized this insight into a six-element instructional model: modeling, coaching, scaffolding, articulation, reflection, and exploration. The model describes precisely how tacit expert knowledge transfers to a novice through embedded practice with feedback — and why the transfer fails when any element is removed.

Ericsson, Krampe, and Tesch-Römer (1993) extended the analysis to expert performance generally, arguing that the active ingredient in expertise development is *deliberate practice with immediate feedback from a teacher or coach*, not unstructured experience or sheer hours invested. While subsequent replication work (Macnamara & Maitra, 2019) has tempered the original effect size estimates, the core finding — that feedback from an external expert is structurally distinct from solo practice — remains well supported.

The most provocative claim in this literature comes from Bloom (1984), whose “2 Sigma Problem” reported that one-to-one tutoring combined with mastery learning produced a two-standard-deviation advantage over conventional group instruction. Subsequent replication work (VanLehn, 2011; Nintil, 2018) has placed the realistic effect size closer to 0.79 standard deviations — still an extraordinarily large effect by social science standards, but smaller than Bloom’s original figure. The phenomenon is real. The magnitude is contested. The relevant point for present purposes is that even the conservative estimate of human-tutor advantage exceeds nearly any other documented educational intervention.

2.4 Implementation Facilitation: The Academic Home of “Done-WITH-You”

The most directly relevant literature for the present argument comes from the field of Implementation Science. The Consolidated Framework for Implementation Research, introduced by Damschroder and colleagues (2009) and updated in 2022 (Damschroder et al., 2022), integrates nineteen prior frameworks into five domains — intervention characteristics, outer setting, inner setting, individual characteristics, and process — and treats *facilitation* as a discrete, named implementation strategy. The Exploration, Preparation, Implementation, Sustainment (EPIS) framework introduced by Aarons, Hurlburt, and Horwitz (2011) goes further, formally identifying “bridging factors” as the connective tissue between an external expert and the internal organization undergoing change.

Powell and colleagues (2015), in the Expert Recommendations for Implementing Change (ERIC) compilation, catalogued seventy-three discrete implementation strategies and consistently ranked facilitation, ongoing consultation, and centralized technical assistance among the highest-endorsed by expert panels. Stetler, Legro, Rycroft-Malone, and colleagues (2006), in a qualitative evaluation of facilitation experiences in the U.S. Veterans Health Administration, defined external facilitation as “a deliberate and valued process of interactive problem solving and support... in the context of a recognized need for improvement and a supportive interpersonal relationship.” This definition is, almost word for word, the operational definition of done-with-you implementation.

Empirical comparisons within this literature consistently favor facilitated over unfacilitated implementation. Kirchner and colleagues (2014), in a quasi-experimental matched-pair study of Primary Care–Mental Health Integration, found measurably higher reach at facilitated sites than at non-facilitated controls. Ritchie, Parker, and Kirchner (2017) found in a qualitative analysis that facilitation fosters higher-quality programs adhering to evidence specifically *in challenged settings* — that is, in contexts where unaided implementation would predictably fail.

The literature on Implementation Facilitation comprises more than five hundred peer-reviewed publications between 1996 and 2021 (U.S. Veterans Health Administration QUERI Implementation Facilitation Literature Collection). To the authors' knowledge, none of this work has been formally translated into the SMB-AI context, nor into the Hispanic SMB context specifically.

2.5 The Hispanic SMB Context

The economic backdrop matters. Stanford's *State of Latino Entrepreneurship* (Orozco et al., 2025; Stanford Graduate School of Business, 2026) has documented over a decade that Latino-owned businesses are growing faster than the national average and increasingly contribute the majority of net new firms in California and Florida. The 2024 edition of the report found that 14% of Latino-owned businesses with revenues exceeding one million dollars were using AI, compared to 7% of comparable white-owned businesses — a notable inversion of the assumed adoption gap. The 2025 edition documented that Latino-owned firms doubled their AI adoption from 2024 to 2025, bringing the population to roughly the same overall adoption rate as white-owned firms (approximately 20%).

McKinsey & Company (2024) estimated that closing the parity gap for U.S. Latino-owned small and medium businesses would unlock approximately 1.4 trillion dollars in additional revenue and create five to six million net new jobs. In Latin America, the picture is more constrained: Olvera-Vera, Peñarreta-Barrera, Alvear-Dávalos, and Yáñez-Escobar (2025), in a comparative study of eleven Latin American countries published in the *Multidisciplinary Latin American Journal*, identified limited connectivity, scarce skilled talent, restricted financing, cultural resistance, and data-security concerns as the dominant structural barriers to digital transformation among regional SMEs.

The combined picture is one of high-velocity adoption among a population that is structurally disadvantaged on the cost side of failed implementation. Latino owner-operators have less access to capital, fewer recoveries from failed experiments, and are statistically more likely to have founded their businesses as foreign-born immigrants navigating a second-language operating environment. The DIY default is, in this population, a particularly expensive default to follow.

2.6 Research Gap

The literature documents, in six independent streams, that expert-mediated learning and implementation outperform self-directed approaches by large effect sizes. The literature also documents, separately, that Hispanic-owned SMBs are adopting AI rapidly while operating with

structural constraints that compound the cost of failure. What the literature does *not* yet contain is a synthesis that translates the well-validated Implementation Facilitation construct into the specific context of AI adoption in Hispanic-owned SMBs, nor a framework that articulates a third path between pure DIY and opaque done-FOR-you outsourcing. The remainder of this paper develops that synthesis.

3. Analysis and Discussion

3.1 The Replicated Two-to-One Pattern

When the evidence assembled in Section 2 is placed side by side, a recurring quantitative pattern emerges. The MIT NANDA findings (Challapally et al., 2025) report a 67% to 33% success ratio between vendor-partnered and internally built generative AI deployments — a two-to-one differential. SCORE (2018) reports that mentored small businesses survive past five years at approximately 70%, compared to roughly 35% for unmentored counterparts — also two-to-one. Userpilot’s SaaS data (2023) shows that products with structured guided onboarding achieve activation rates approximately 50% higher than self-service equivalents, reaching the 60–70% activation range against the 37.5% self-service average. Bloom’s tutoring effect, even in conservative replications (VanLehn, 2011), places one-to-one expert guidance at roughly 0.79 standard deviations above unguided learning — a magnitude that, in the social sciences, almost no other intervention matches.

These figures are drawn from different populations, different time periods, different methodologies, and different outcome variables. That they converge on a similar order of magnitude is meaningful. A pattern that replicates across coaching meta-analyses, mentoring cohort data, SaaS activation telemetry, enterprise AI deployments, and laboratory studies of human tutoring is, in all likelihood, capturing a structural property of human skill acquisition rather than a context-specific quirk.

The structural property is straightforward: when a person attempts to apply an unfamiliar tool to an unfamiliar problem, the rate-limiting variable is rarely access to the tool. It is the presence or absence of someone who has previously navigated the same application and can intervene at the moment of friction. The MIT NANDA report names this absence as the “learning gap”; Implementation Science names it as the absence of “facilitation”; Vygotsky names it as the unfilled Zone of Proximal Development. The mechanism is the same.

3.2 Why DIY Appears Rational Despite the Evidence

If the empirical record is this consistent, why does the DIY default persist? The answer lies in a set of well-documented cognitive biases that systematically distort the cost-benefit calculus for owner-operators.

The first is what Maslow (1966) named the Law of the Instrument and what behavioral economists subsequently formalized as the *saliency bias*: when a powerful new tool becomes

available, attention concentrates disproportionately on the tool itself rather than on the structural problems for which the tool may or may not be appropriate. The owner-operator who has just discovered ChatGPT applies it to the most visible task — drafting emails, generating social posts, summarizing documents — rather than the most impactful task, which is usually invisible by definition because it sits at the structural level of how the business is designed.

The second is the *knowing-doing gap* (Pfeffer & Sutton, 2000), which describes the well-documented phenomenon in organizational behavior whereby information about what works fails to translate into action. This is precisely the mechanism documented by the MOOC completion data: registrants who paid for certificates and signaled clear intent nonetheless dropped out at rates exceeding 50%. The information was available. The action did not follow.

The third is *cost visibility asymmetry*. The fee paid to an external implementer is concrete, line-itemed, and presented for decision in a single moment. The cost of failed self-implementation is diffuse: a month spent watching tutorials, an automation that broke after two weeks, a competitor who pulled ahead, a family vacation cancelled because the owner could not step away. Each of these costs is real and, in aggregate, substantial. But none of them generate an invoice. The cognitive system that evaluates the consultant fee never sees the comparable cost on the other side of the ledger.

The fourth is *choice overload* (Iyengar & Lepper, 2000), which describes the documented decrease in decision quality as the number of available options increases. The contemporary AI tools market presents the owner-operator with hundreds of platforms, dozens of training courses, thousands of prompt-engineering tutorials, and a steady stream of competing claims about what to build first. Choice overload converts the rational process of “select the right tool for the right problem” into the documented psychological state of decision paralysis — at which point the owner-operator either disengages entirely or selects on the basis of the most superficial cue, typically price.

These four biases compound. Together they explain why the DIY default is selected even by capable, well-informed owners who have access to the same data presented above. The biases are not signs of poor judgment. They are predictable failures of an unaided cognitive system operating in a high-noise environment.

3.3 The Cost of Solving the Wrong Problem

A consequence of the above biases is that self-directed adopters tend to apply new tools to the most visible task rather than the most impactful one. Consider, by way of illustration, a contractor who notices a competitor has installed a chatbot on their website and concludes that they should do the same. The chatbot is purchased, installed, and configured. It functions technically. But the contractor’s actual revenue leak — discovered only when an external party with sales judgment examines the operation — was that 40% of inbound calls were going unanswered because there was no immediate-response system, and the quotes that were sent followed no commercial structure capable of generating purchase confidence. The chatbot was the right an-

swer to the wrong question. It cost money. It produced no revenue lift. And the underlying structural pattern — a business in which every revenue-generating activity ran through a single person — remained invisible.

The economic magnitude of this misapplication is non-trivial. If a contractor receives twenty inbound leads per month at an average ticket of five thousand dollars and loses 40% to slow response, the monthly opportunity cost is forty thousand dollars. Over a year, that is nearly half a million dollars. A chatbot at fifty dollars per month does not address that loss. A facilitated implementation that diagnoses the actual bottleneck — beginning with response infrastructure and quote structure rather than presence on the website — does. The differential is not the cost of the tool. It is the presence or absence of judgment about which tool, applied to which problem, in which order.

This observation generalizes. The distinguishing characteristic of expert-mediated implementation is not access to better technology; the underlying GPT, Claude, or Gemini instance is essentially identical between guided and unguided users. The distinguishing characteristic is the diagnostic process that precedes tool selection. The tool is the vehicle. Judgment is what determines whether the vehicle moves toward revenue or away from it.

3.4 Proposed Framework: The *Agentes Para Tu Negocio* Model

The *Agentes Para Tu Negocio* model proposes a structurally distinct middle path between two failing alternatives. Pure DIY adoption produces the documented two-to-one disadvantage in outcome. Pure done-FOR-you outsourcing — the traditional consultancy model — succeeds operationally but creates a different failure mode: the client receives a working system but does not develop the judgment to maintain, extend, or correctly diagnose its limits. When the consultant departs, the system degrades. The dependency is merely transferred.

The proposed alternative is *done-WITH-you* implementation, situated theoretically within the Implementation Facilitation construct established by Stetler, Kirchner, and colleagues. The model has four constitutive elements, each derived from the literature reviewed above.

The first is *bottleneck-first diagnosis*. Rather than beginning with the question “which tool would you like installed?” — the question that produces solution-first bias — the engagement begins with structural diagnosis: where is revenue actually leaking, where is the owner’s time actually consumed, and what single intervention would produce the highest immediate return. This corresponds directly to the Exploration phase of the EPIS framework (Aarons et al., 2011) and to the modeling and articulation elements of cognitive apprenticeship (Collins et al., 1989).

The second is *codification of tacit business judgment*. The work of the implementer is not to install a generic tool with the client’s data, but to encode the specific commercial logic of the client’s business — pricing structure, follow-up cadence, qualification criteria, brand voice — into the system being built. This is the operational form of Polanyi’s (1966) insight that the most valuable knowledge in any expert practice is tacit and cannot be transferred through documentation alone. It requires demonstration in context.

The third is *embedded learning during delivery*. The client is not handed a finished system at the end of the engagement; the client observes the system being built and develops, as a byproduct, a mental model of how the system works and what it reveals about the structure of their own business. This corresponds to the *articulation* and *reflection* elements of the cognitive apprenticeship model (Collins et al., 1989) and is the structural answer to the dependency problem of pure done-FOR-you outsourcing. The dependency is not transferred from owner to consultant; it is dissolved by the development of judgment in the owner.

The fourth is *first principles, then expansion*. The initial engagement is deliberately bounded — one bottleneck, one system, one delivery — rather than a multi-month transformation program. This corresponds to the implementation science principle of *bounded scope as a precondition for sustainability* (Powell et al., 2015) and addresses the cognitive overload that occurs when an owner-operator faces an unbounded change initiative on top of existing operational demands.

The model is not a product configuration. It is a structural argument: the active ingredient in successful AI implementation in SMBs is the combination of business judgment, tacit knowledge codification, embedded teaching, and bounded scope. The technology — whichever generative AI platform happens to be current — is the vehicle. The four elements above are what determine whether the vehicle moves the business or merely consumes the budget.

3.5 Practical Implications

For the owner-operator, the practical implication is that the question “should I do this myself or hire someone?” is ill-formed. The relevant question is whether the implementation pathway being considered includes (a) diagnosis before tool selection, (b) codification of one’s own commercial judgment into the system, (c) opportunity to develop a working mental model during delivery, and (d) bounded initial scope. If a self-directed approach can incorporate all four — through a combination of reading, mentorship, peer review, and disciplined sequencing — the differential narrows. If a paid engagement omits any of them, the cost may not be recovered. The fee is not the relevant variable. The structural integrity of the implementation pathway is.

For policymakers and ecosystem actors serving Hispanic-owned SMBs, the practical implication is that the dominant intervention model — free or low-cost AI training programs offered as self-paced courses — is structurally mismatched to the cliff of disengagement documented in Section 2.2. Resources that would produce higher returns are those that fund facilitated implementation pathways, particularly in Spanish, with implementers who carry both AI fluency and business operating experience. The McKinsey (2024) estimate of 1.4 trillion dollars in unrealized parity-gap value provides a magnitude argument for such investment.

For practitioners offering implementation services, the practical implication is that the marketing language used by most AI consultancies — “personalized chatbots,” “custom AI solutions,” “intelligent automation” — describes the vehicle and not the value, and consequently commodifies the offering. A practitioner who genuinely operates a done-with-you model has a positioning problem to solve: how to communicate, in a market saturated with vehicle-language,

that what is being offered is structural diagnosis and embedded teaching. That is a substantive marketing question, and the answer lies in case-based demonstration rather than feature-based description.

4. Conclusions

4.1 Summary of Findings

The empirical record across six independent research streams supports a single, robust conclusion: expert-mediated implementation outperforms self-directed implementation on measurable outcomes, with effect sizes consistent enough across domains and time periods to constitute a structural finding rather than a context-specific result. The mechanism is the presence of expert judgment during the critical application phase of implementation — what the Implementation Science literature calls *facilitation* and what the present paper has referred to throughout as *criterio*. The mechanism is not the tool. Two implementers using the same generative AI platform produce divergent business outcomes when one carries judgment about what to build first and the other does not. The tool is the vehicle. Judgment is the active ingredient.

For Hispanic-owned SMBs specifically — a population whose AI adoption rate doubled between 2024 and 2025, whose access to capital remains structurally constrained, and whose Latin American counterparts face additional barriers of connectivity, talent, and language — the cost of selecting the DIY default on the basis of cost-visibility bias is particularly steep. The proposed *Agentes Para Tu Negocio* model, situated within the Implementation Facilitation construct of established implementation science, offers a structurally appropriate third path that neither leaves the owner alone with an unfamiliar tool nor transfers dependency to an opaque outsourcer.

4.2 Limitations

This review is subject to the selection bias inherent in narrative literature reviews. The convergence of effect sizes across domains is suggestive but does not constitute a controlled experiment of the proposed framework. The 2025 MIT NANDA report, while methodologically diverse, has not yet been peer-reviewed, and several of the most-cited industry figures — the Manchester coaching ROI, the SCORE survival differential, the SaaS retention curves — rely on perceptual self-report or industry-commissioned data rather than randomized designs. The widely circulated “70% of digital transformations fail” statistic, traced to Hammer and Champy’s (1993) original “unscientific estimate,” has been deliberately excluded from the analysis above on grounds of methodological insufficiency, though it is acknowledged here as a data point in the broader narrative.

The proposed *Agentes Para Tu Negocio* framework has not yet been validated through controlled experimental studies. Its constitutive elements are each independently grounded in the cited literature, but the package as deployed in the field remains, as of this writing, in the conceptual and case-based phase of evidence development.

4.3 Future Research Directions

Three lines of future research follow naturally. The first is a controlled comparative study within Hispanic-owned SMBs randomizing access to facilitated versus self-directed AI implementation pathways, with measurable outcomes on revenue, owner time recovered, and system durability over twelve to twenty-four months. The second is a longitudinal case-series of implementations conducted under the *Agentes Para Tu Negocio* model, documenting the specific bottlenecks diagnosed, the systems built, and the longitudinal performance of those systems against the implementation science success criteria established by Damschroder and colleagues (2022). The third is an applied translation of the EPIS framework's "bridging factors" construct into a vocabulary appropriate for the SMB-AI context, in both English and Spanish, with the explicit goal of equipping owner-operators to evaluate prospective implementation engagements against criteria stronger than price.

The broader research program of which this paper forms one component will examine related questions in adjacent domains: the comparative outcomes of synchronous versus asynchronous expert mediation; the role of language and cultural context as moderators of implementation success in Latino-owned firms; and the structural features that distinguish implementation engagements that produce durable systems from those that produce visible activity without durable change. The intended trajectory is the development of an evidence-based framework specific to the intersection of AI adoption, owner-operated SMBs, and the Hispanic market — a synthesis that the existing literature has not yet produced.

Resumen en Español

La narrativa dominante sobre la adopción de inteligencia artificial generativa en pequeñas y medianas empresas (pymes) trata a la herramienta misma como el determinante principal del resultado. Bajo este encuadre, el dueño-operador racional concluye que contratar a un implementador externo es un gasto evitable: un costo visible a través de una factura, contrastado contra herramientas que parecen gratuitas o casi gratuitas. Este artículo examina si esa conclusión resiste el escrutinio empírico. Sintetizando evidencia de seis líneas de investigación independientes — meta-análisis de coaching ejecutivo, datos de deserción en cursos masivos en línea, métricas de incorporación en SaaS, literatura de Ciencia de la Implementación sobre facilitación, investigación reciente sobre despliegue de IA generativa empresarial, y psicología educativa sobre andamiaje y práctica deliberada — el análisis revela un patrón consistente y replicado: la implementación mediada por expertos supera a la adopción autodirigida en proporciones cercanas a dos a uno en los resultados medibles. El reporte 2025 del proyecto MIT NANDA encuentra que los despliegues de IA con socios externos tienen éxito aproximadamente el 67% de las veces, frente al 33% para construcciones internas autodirigidas. Proporciones comparables aparecen en datos de mentoría a pequeñas empresas de SCORE, en métricas de activación de SaaS, y en décadas de investigación sobre facilitación en servicios de salud. El artículo introduce el modelo *Agentes Para Tu Negocio* — un enfoque "hecho-CONTIGO" enmarcado dentro

del constructo de Facilitación de la Implementación establecido por Stetler, Kirchner y colegas — como una respuesta estructuralmente apropiada para pymes hispanas, una población cuya tasa de adopción de IA se duplicó entre 2024 y 2025 pero permanece desatendida por la falsa dicotomía entre el “hazlo tú solo” puro y la subcontratación opaca de “hecho-POR-ti”. El modelo propuesto trata el criterio de negocio, y no la herramienta, como el ingrediente activo.

Palabras clave: implementación guiada, adopción tecnológica mediada por expertos, ROI de coaching, ciencia de la implementación, autodidacta vs asistido, consultoría para pymes, eficacia del onboarding, Agentes Para Tu Negocio, adopción de IA en negocios hispanos, facilitación de implementación, IA generativa pymes

References

- Aarons, G. A., Hurlburt, M., & Horwitz, S. M. (2011). Advancing a conceptual model of evidence-based practice implementation in public service sectors. *Administration and Policy in Mental Health and Mental Health Services Research, 38*(1), 4–23. <https://doi.org/10.1007/s10488-010-0327-7>
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4–16. <https://doi.org/10.3102/0013189X013006004>
- Challapally, A., et al. (2025). *The GenAI divide: State of AI in business 2025*. MIT NANDA Project, Massachusetts Institute of Technology.
- Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453–494). Lawrence Erlbaum Associates.
- Damschroder, L. J., Aron, D. C., Keith, R. E., Kirsh, S. R., Alexander, J. A., & Lowery, J. C. (2009). Fostering implementation of health services research findings into practice: A consolidated framework for advancing implementation science. *Implementation Science, 4*, 50. <https://doi.org/10.1186/1748-5908-4-50>
- Damschroder, L. J., Reardon, C. M., Widerquist, M. A. O., & Lowery, J. C. (2022). The updated Consolidated Framework for Implementation Research based on user feedback. *Implementation Science, 17*, 75. <https://doi.org/10.1186/s13012-022-01245-0>
- Ely, K., Boyce, L. A., Nelson, J. K., Zaccaro, S. J., Hernez-Broome, G., & Whyman, W. (2010). Evaluating the effectiveness of executive coaching: Beyond ROI? *Coaching: An International Journal of Theory, Research and Practice, 3*(2), 91–110. <https://doi.org/10.1080/17521882.2010.493213>
- 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>
- Hammer, M., & Champy, J. (1993). *Reengineering the corporation: A manifesto for business revolution*. HarperBusiness.

- International Coaching Federation. (2023). *2023 ICF global coaching study: Executive summary*. International Coaching Federation. <https://coachingfederation.org/wp-content/uploads/2023/04/2023ICFGlobalCoachingStudyExecutiveSummary.pdf>
- International Coaching Federation. (2025). *2025 ICF global coaching study: Executive summary*. International Coaching Federation. <https://coachingfederation.org/resource/2025-icf-global-coaching-study-executive-summary/>
- Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6), 995–1006. <https://doi.org/10.1037/0022-3514.79.6.995>
- Jones, R. J., Woods, S. A., & Guillaume, Y. R. F. (2016). The effectiveness of workplace coaching: A meta-analysis of learning and performance outcomes from coaching. *Journal of Occupational and Organizational Psychology*, 89(2), 249–277. <https://doi.org/10.1111/joop.12119>
- Jordan, K. (2014). Initial trends in enrolment and completion of massive open online courses. *International Review of Research in Open and Distance Learning*, 15(1), 133–160. <https://doi.org/10.19173/irrodl.v15i1.1651>
- Jordan, K. (2015). Massive open online course completion rates revisited: Assessment, length and attrition. *International Review of Research in Open and Distributed Learning*, 16(3), 341–358. <https://doi.org/10.19173/irrodl.v16i3.2112>
- Kirchner, J. E., Ritchie, M. J., Pitcock, J. A., Parker, L. E., Curran, G. M., & Fortney, J. C. (2014). Outcomes of a partnered facilitation strategy to implement Primary Care–Mental Health. *Journal of General Internal Medicine*, 29(Suppl 4), S904–S912. <https://doi.org/10.1007/s11606-014-3027-2>
- Macnamara, B. N., & Maitra, M. (2019). The role of deliberate practice in expert performance: Revisiting Ericsson, Krampe & Tesch-Römer (1993). *Royal Society Open Science*, 6(8), 190327. <https://doi.org/10.1098/rsos.190327>
- Maslow, A. H. (1966). *The psychology of science: A reconnaissance*. Harper & Row.
- McGovern, J., Lindemann, M., Vergara, M., Murphy, S., Barker, L., & Warrenfeltz, R. (2001). Maximizing the impact of executive coaching: Behavioral change, organizational outcomes, and return on investment. *The Manchester Review*, 6(1), 1–9.
- McKinsey & Company. (2024). *The economic state of Latinos in America: Building up small businesses*. McKinsey & Company. <https://www.mckinsey.com/featured-insights/diversity-and-inclusion/the-economic-state-of-latinos-in-the-us>
- Olvera-Vera, J., Peñarreta-Barrera, M., Alvear-Dávalos, F., & Yáñez-Escobar, P. (2025). Transformación digital de las PYMES en América Latina: Barreras, oportunidades y estrategias para la competitividad. *Multidisciplinary Latin American Journal*, 3(2). <https://doi.org/10.62452/mlaj.v3i2.98>
- Orozco, M., Chávez Zárata, R., & Foster, G. (2025). *State of Latino entrepreneurship 2024*. Stanford Latino Entrepreneurship Initiative, Stanford Graduate School of Business.

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*. Doubleday.

Powell, B. J., Waltz, T. J., Chinman, M. J., Damschroder, L. J., Smith, J. L., Matthieu, M. M., Proctor, E. K., & Kirchner, J. E. (2015). A refined compilation of implementation strategies: Results from the Expert Recommendations for Implementing Change (ERIC) project. *Implementation Science, 10*, 21. <https://doi.org/10.1186/s13012-015-0209-1>

Reich, J., & Ruipérez-Valiente, J. A. (2019). The MOOC pivot. *Science, 363*(6423), 130–131. <https://doi.org/10.1126/science.aav7958>

Ritchie, M. J., Parker, L. E., & Kirchner, J. E. (2017). Using implementation facilitation to foster clinical practice quality and adherence to evidence in challenged settings: A qualitative study. *BMC Health Services Research, 17*, 294. <https://doi.org/10.1186/s12913-017-2217-0>

SCORE. (2018). *The mentoring effect on small business survival and growth*. SCORE Association.

Stanford Graduate School of Business. (2026). *State of Latino entrepreneurship 2025: Eleventh annual report*. Stanford Latino Entrepreneurship Initiative.

Stetler, C. B., Legro, M. W., Rycroft-Malone, J., Bowman, C., Curran, G., Guihan, M., Hagedorn, H., Pinos, S., & Wallace, C. M. (2006). Role of “external facilitation” in implementation of research findings: A qualitative evaluation of facilitation experiences in the Veterans Health Administration. *Implementation Science, 1*, 23. <https://doi.org/10.1186/1748-5908-1-23>

Theeboom, T., Beersma, B., & van Vianen, A. E. M. (2014). Does coaching work? A meta-analysis on the effects of coaching on individual level outcomes in an organizational context. *The Journal of Positive Psychology, 9*(1), 1–18. <https://doi.org/10.1080/17439760.2013.837499>

Userpilot. (2023). *State of SaaS onboarding 2023*. Userpilot. <https://pages.userpilot.com/state-of-saas-2023/>

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist, 46*(4), 197–221. <https://doi.org/10.1080/00461520.2011.61>

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.