

The Information vs. Implementation Divide: How ChatGPT Exposes the Structural Weakness of Traditional Online Courses

Humberto Inciarte

May 2026

Abstract

A widely circulated narrative in the creator economy holds that ChatGPT and consumer-facing large language models are killing the online course market. Anecdotal evidence — declining sales for individual creators, the public collapse of information-services firms such as Chegg, and surveys showing that learners increasingly use ChatGPT to substitute for paid educational content — has reinforced this view. This paper argues that the prevailing narrative is empirically incomplete and structurally misleading. Drawing on peer-reviewed studies in educational research, behavioral psychology, knowledge management, and the labor economics of generative AI, the paper documents that (a) the completion crisis in self-paced online courses substantially predates ChatGPT, (b) the commoditization of explicit information was already predicted by the economics of information goods, and (c) the post-ChatGPT competitive frontier in adult education lies in the implementation gap — the documented failure of information transfer alone to produce behavior change and skill acquisition. Market evidence from creator-economy platforms whose value proposition centers on implementation and community, rather than on information delivery, supports this thesis: such platforms have continued to grow during the same period in which information-only services have contracted. The paper proposes the *Information–Implementation Divide* as an integrative framework and introduces the CursoVivo implementation model — a methodology-encoded artificial intelligence layer designed to operationalize personalized implementation at scale within an existing course — as a concrete instantiation of the framework. Implications for course creators, EdTech researchers, and the design of post-LLM digital learning systems are discussed.

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

ORCID: 0009-0006-0379-7328

Corresponding author: hola@humbertoinciarte.com

Keywords: ChatGPT online courses; AI impact education; e-learning disruption; course creator strategy; information commodity; implementation gap; online education AI; CursoVivo; completion rates; tacit knowledge; MOOC; generative AI in education; post-LLM learning design.

1. Introduction

1.1 The Prevailing Narrative

When OpenAI released ChatGPT on November 30, 2022, it reached an estimated 100 million monthly active users within approximately two months. A widely cited UBS analyst note characterized the trajectory as the fastest consumer-internet adoption ramp in twenty years of analyst coverage (Hu, 2023). By early 2024, the Stanford Institute for Human-Centered AI reported that organizational adoption of generative AI had risen from 33% to 71% in a single year, while the cost of querying a model of GPT-3.5-equivalent capability had fallen by approximately 280-fold between November 2022 and October 2024 (Stanford HAI,

2024). For the first time in the history of digital education, learners had near-zero-marginal-cost access to a system that could explain, summarize, and contextualize most of the explicit knowledge previously sold inside paid online courses.

Against this backdrop, a now-familiar narrative crystallized in the creator economy: *ChatGPT is killing online courses*. Course creators reported declining lifetime conversion rates, prospects citing free alternatives during sales conversations, and a generalized sense that the value proposition of recorded information-delivery courses had been hollowed out. Public-market evidence reinforced the narrative. On May 2, 2023, the education-services firm Chegg lost approximately half of its market capitalization in a single trading session after acknowledging on its earnings call that ChatGPT was materially affecting new-customer acquisition (Sherman, 2023). The headline was repeated across financial and trade press: artificial intelligence was, in this telling, an extinction-level event for paid digital education.

This narrative has clear surface validity. Surveys conducted within months of ChatGPT's release found that the majority of medical and research trainees had already integrated the tool into their learning workflows, often substituting it for textbook explanations and paid coursework (Hosseini et al., 2023). Empirical testing of ChatGPT against the question banks of major massive open online course (MOOC) platforms showed that the model could correctly answer a substantial portion of the assessment material, raising direct substitution concerns (Acuña Caicedo et al., 2023). It is therefore unsurprising that a defensive response has predominated among creators: lower prices, increase advertising spend, produce more content, or exit the market.

1.2 The Problem

The defensive response is costly and, this paper argues, structurally misdirected. Reducing prices to compete with a free, near-infinite source of generic information is unlikely to yield favorable economics. Increasing advertising spend without addressing what happens *after* the sale compounds an underlying weakness rather than resolving it. Producing additional content multiplies the volume of material that learners must traverse without improving the probability that any of it will be applied. More fundamentally, the prevailing narrative treats online courses as a homogeneous category, when in fact the empirical literature suggests that the courses most exposed to ChatGPT and the courses least exposed to it differ along a structural dimension that has been visible — but largely unaddressed — in the educational research literature for over a decade.

That dimension is the distinction between *information* and *implementation*: between the transfer of explicit content and the production of measurable behavior change in the learner. The hypothesis advanced in this paper is that ChatGPT has not introduced a new disruption to the online education market. It has *exposed* a structural weakness that always existed in courses whose entire value proposition reduced to information transfer, while leaving largely intact — and in some segments accelerating — the value of courses, programs, and learning systems whose proposition is implementation, transformation, and personalized application.

1.3 Research Question and Thesis

This paper examines the following question: *Does the empirical literature support the claim that ChatGPT is killing the online course market, or does it support a more granular interpretation in which the commoditization of explicit information is exposing a pre-existing structural weakness in information-only courses?*

The thesis advanced is the latter. Specifically, the paper argues that:

1. The completion crisis in self-paced online education substantially predates ChatGPT and reflects a structural feature of information-only delivery, not a recent disruption.
2. The commoditization of explicit information was already predicted by the economics of information goods more than two decades before generative AI reached consumer scale.
3. The post-ChatGPT competitive frontier in adult education lies in the *implementation gap* — the documented failure of information transfer alone to produce behavior change — and in the externalization of proprietary, tacit methodology that cannot be replicated by general-purpose models.

The paper proposes the *Information–Implementation Divide* as an integrative framework and introduces the *CursoVivo* implementation model as a concrete instantiation. The argument is developed across four parts: a literature review (§2) synthesizing four siloed bodies of research; an analysis of post-ChatGPT market evidence and a structural account of the divide (§3); the proposed framework and its practical implications; and concluding remarks on limitations and future research (§4).

2. Literature Review

2.0 Methodological Note

This review synthesizes peer-reviewed empirical studies, working papers from established research institutions, and institutional reports published primarily between 2019 and 2025, with foundational works extending to 1984. Sources were drawn from PubMed, Google Scholar, the National Bureau of Economic Research (NBER), arXiv, Stanford HAI, and McKinsey Global Institute publications. Inclusion criteria prioritized studies with measurable learning, productivity, or market outcomes relevant to digital adult education. Industry data from creator-economy platforms (Hotmart, Kajabi) were included where peer-reviewed evidence was unavailable, with primary corporate communications cited rather than secondary aggregators. Where the popular narrative cited statistics that could not be traced to a primary source, those statistics were excluded from the analysis.

2.1 The Pre-existing Structural Weakness of Online Courses

The most authoritative empirical evidence on completion rates in self-paced online education comes from the multi-year MITx and HarvardX program. In a *Science* paper published nearly four years before ChatGPT's release, Reich and Ruipérez-Valiente (2019) analyzed all 565 MOOCs offered by MIT and Harvard between 2012 and 2018. Their central finding directly contradicts the post-ChatGPT framing: completion rates had not improved over six years of platform optimization, the vast majority of learners did not return after their first year, and growth in MOOC participation had concentrated in the world's most affluent

countries — a population already saturated with educational opportunities. The structural weakness, in other words, was demonstrably a property of the delivery model, not a consequence of competition from generative AI.

This finding aligns with a broader pattern in the educational research literature. Bloom (1984), in a seminal article in *Educational Researcher*, documented what has come to be known as the “2 sigma problem”: students receiving one-to-one tutoring with mastery learning performed approximately two standard deviations above students in conventional classroom instruction, with the average tutored student outperforming 98% of the comparison group. Bloom framed the central challenge of educational design as identifying methods of group instruction that could approximate the effectiveness of individualized tutoring. The challenge has remained largely unsolved at scale; static-video MOOCs, despite their reach, have not narrowed the personalization gap that Bloom identified more than four decades ago.

Practitioner data are consistent with this academic picture. Cohort-based and high-touch online programs that combine information delivery with structured implementation, accountability, and feedback report completion rates approaching one order of magnitude above MOOC averages — Harvard Business School Online has reported completion rates near 85%, and high-engagement programs such as Seth Godin’s altMBA have reported completion rates near 96% (Habif, 2017). The variable is not the population of learners; it is the structure of the delivery.

2.2 Information Goods Economics and the Commoditization of Knowledge

That generic information would eventually be commoditized was not an unanticipated consequence of generative AI; it was a prediction of the economics of information goods made well before consumer-facing large language models existed. Shapiro and Varian (1999), in *Information Rules*, argued that the marginal cost of reproducing digital information goods approaches zero, that the equilibrium price of undifferentiated information therefore tends toward zero in competitive markets, and that sustainable differentiation must come from one of several alternative axes: versioning, lock-in, network effects, brand, or — most relevantly for the present paper — complementary services that the information itself cannot provide.

Generative AI did not invent this dynamic. It accelerated it. The Stanford AI Index documented that the cost of querying a GPT-3.5-equivalent model fell approximately 280-fold between late 2022 and late 2024 (Stanford HAI, 2024), a collapse that effectively brought the marginal cost of accessing high-quality explicit information to near zero for the average end user. Brynjolfsson, Li, and Raymond (2025), in a field experiment with 5,172 customer-support agents published in *The Quarterly Journal of Economics*, found that generative AI assistance increased productivity by approximately 15% on average, with effects concentrated in novice and lower-skilled workers (gains exceeding 30% in some cases) and approaching zero for experienced experts. The authors interpret this asymmetric pattern in terms that are directly relevant to the present argument: AI systems “may be capable of capturing and disseminating the behaviors of the most productive agents,” functioning as an amplification mechanism for expert tacit knowledge rather than as a substitute for it (Brynjolfsson et al., 2025, p. 891).

The implication for the online course market is straightforward and consistent with Shapiro and Varian’s

earlier framework. The layer of paid online education that consisted of explicit, codifiable information — the layer most analogous to the support scripts and FAQ knowledge bases that ChatGPT readily reproduces — was the layer most exposed to commoditization. The layer that consists of personalized application of proprietary methodology, accountability structures, and feedback-rich practice is, by the same economic logic, the layer most insulated from it.

2.3 The Knowing-Doing Gap and Implementation Science

A separate body of literature, developed in parallel to the educational research tradition, addresses why information transfer alone systematically fails to produce behavior change. Pfeffer and Sutton (2000), in *The Knowing-Doing Gap*, documented across multiple organizational case studies that improvements in performance “depend largely on implementing what is already known, rather than on adopting new or previously unknown ways of doing things” (p. ix). The bottleneck, in other words, is rarely the absence of relevant knowledge; it is the absence of the structures, incentives, and mechanisms required to act on knowledge that is already available.

Behavioral psychology has produced a substantial body of converging evidence. Gollwitzer (1999), introducing the concept of *implementation intentions*, demonstrated that specifying when, where, and how one will act on a goal — through if-then plans linking situational cues to specific actions — substantially increases rates of goal-directed behavior compared with goal intentions alone. A subsequent meta-analysis of 94 independent studies found a medium-to-large effect of implementation intentions on goal attainment across domains (Gollwitzer & Sheeran, 2006). Ericsson, Krampe, and Tesch-Römer (1993), in their foundational work on deliberate practice, established that expert performance is the product of feedback-rich, goal-directed activity sustained over years — not of passive exposure to information about the domain. Roediger and Karpicke (2006) demonstrated that retrieval practice — active recall through low-stakes testing — produces substantially greater long-term retention than re-reading or passive review, despite learners systematically predicting the opposite.

Each of these literatures was developed independently of any AI context. Together they describe a population of mechanisms — implementation intentions, deliberate practice with feedback, retrieval practice, structured accountability — that produce learning outcomes and that none of them depend on information access being scarce. None of them are delivered, in any meaningful form, by a static video library; and none of them are delivered, in any meaningful form, by a general-purpose conversational AI that has no model of the learner’s current state, no record of prior commitments, and no methodology-specific lens through which to interpret the learner’s situation.

2.4 Tacit vs. Explicit Knowledge in the AI Era

The distinction between information that can be straightforwardly commoditized and methodology that cannot has a longer pedigree in knowledge-management research. Nonaka and Takeuchi (1995) distinguished between *explicit* knowledge — codifiable in language, manuals, formulas, and procedures — and *tacit* knowledge, which is embodied, experiential, and learned primarily through socialization, apprenticeship, and indirect communication. Their framework argued that competitive advantage in knowledge-intensive organizations flows from the cyclical conversion of tacit into explicit knowledge and back again,

a process they formalized as the SECI model (socialization, externalization, combination, internalization).

The relevance of this distinction to post-ChatGPT online education is direct. Information available on the open web — the substrate on which general-purpose large language models are trained — is, by definition, the explicit layer of knowledge in any given domain. The proprietary methodology of an experienced practitioner — the diagnostic shortcuts, the order in which interventions are sequenced, the recognition of edge cases that distinguish expert from novice judgment — typically lives in the tacit layer until and unless it is systematically externalized through some structured process. A generic conversational AI cannot reproduce what has not been externalized, because what has not been externalized is, by construction, not available in its training data.

The technical literature on retrieval-augmented generation provides empirical support for this conceptual point. Bechard and Marquez Ayala (2024), in research from ServiceNow, documented that generic large-language-model outputs hallucinate at rates incompatible with enterprise or educational use, and that only domain-grounded systems trained on verified knowledge bases meaningfully reduce hallucination. More recent work has shown that even state-of-the-art retrieval-augmented systems struggle to exceed 80% factual accuracy on benchmark tasks without explicit grounding architectures (Wood et al., 2024). The research and practitioner literatures on AI in education have begun to converge on a similar conclusion: the highest-impact educational deployments of generative AI involve carefully scaffolded, methodology-grounded multi-agent systems rather than direct learner access to a general-purpose model. Mollick and colleagues (2024), in work from the Wharton Generative AI Lab, describe the use of AI agent systems to deliver personalized practice simulations with mentors, role-players, and instructor-facing evaluators — the operationalization, in their framing, of Bloom’s 2-sigma problem at scale.

2.5 Research Gap

The four bodies of literature reviewed above — the empirical study of MOOC completion, the economics of information goods, the implementation and behavior-change literature, and the tacit-explicit knowledge framework, together with the recent labor economics of generative AI — have largely been developed in isolation from each other. Educational researchers studying MOOC completion have not, in general, framed their findings in terms of information-goods commoditization. Economists studying generative AI productivity effects have not, in general, drawn the connection to behavior-change literatures developed in social psychology. Practitioners in the creator economy have, to date, lacked an integrative framework that synthesizes these traditions and translates them into design principles for post-LLM online education.

This paper contributes to the literature by proposing such a synthesis. The *Information–Implementation Divide* — articulated in §3 — is offered as an integrative framework that connects the four literatures and yields concrete predictions about which segments of the online course market are exposed to AI-driven commoditization and which are insulated from it. The CursoVivo implementation model is offered as a concrete instantiation of the framework, suitable for design and evaluation in subsequent empirical work.

3. Analysis and Discussion

3.1 Empirical Evidence: What Is Actually Happening Post-ChatGPT

The narrative that “ChatGPT is killing online courses” rests on a small number of high-visibility cases drawn from a single segment of the digital education market. The most prominent — Chegg’s market-capitalization collapse following its May 2023 earnings call (Sherman, 2023) — concerns a firm whose value proposition was substantially the on-demand provision of explicit information: textbook-question solutions, expert-supplied explanations of standardized course material, and tutoring services structured around standardized curricula. Chegg’s exposure to ChatGPT, in the language developed in §2.2, was precisely the exposure that information-goods economics predicts: a firm whose differentiation rested on access to commoditizable explicit information should expect to lose that differentiation when the marginal cost of that information collapses. The Chegg case is therefore best understood not as a leading indicator for the entire online education market, but as a clarifying example of what happens to firms positioned exclusively in the information layer.

The same period during which information-services firms have contracted has seen continued growth in segments of the creator economy whose value propositions are positioned in the implementation layer. Hotmart Company, the parent organization of Teachable since 2023, announced in March 2024 that creators on its platform had collectively generated more than \$10 billion in cumulative earnings, with more than 200,000 active creators serving more than 21 million buyers globally (Hotmart Company, 2024). Kajabi, a competitor platform whose product positioning emphasizes creator-owned course delivery, community, and coaching tools, announced in August 2025 that it had crossed \$10 billion in cumulative creator earnings, having added the most recent billion in approximately five months — a pace consistent with sustained growth rather than contraction (Kajabi, 2025). Goldman Sachs Research has projected that the broader creator economy will roughly double from approximately \$250 billion to \$480 billion by 2027 (Sheridan, 2023).

These data are not incompatible with the narrative that ChatGPT is exerting downward pressure on commodity-information products. They are incompatible with the narrative that ChatGPT is killing the online course market in aggregate. The asymmetric pattern — contraction in firms positioned at the information layer, continued growth at the implementation and community layer — is exactly what the economics of information goods (Shapiro & Varian, 1999) and the labor economics of generative AI (Brynjolfsson et al., 2025) jointly predict.

3.2 The Information–Implementation Divide: A Structural Analysis

The pattern documented in §3.1 admits a structural explanation that integrates the literature reviewed in §2. The post-ChatGPT online education market can be analyzed in terms of three layers:

Layer 1 — Explicit information. The codifiable, web-available portion of any domain’s knowledge base. This layer consists of facts, definitions, procedural descriptions, conceptual explanations, and worked examples that are reproducible from publicly accessible sources. The marginal cost of accessing this layer collapsed sharply with the consumer release of generative AI (Stanford HAI, 2024). Courses whose

value proposition reduces to this layer are subject to the commoditization dynamics described by Shapiro and Varian (1999) and bear the highest exposure to AI-driven displacement.

Layer 2 — Proprietary methodology. The systematized, structured judgment of an experienced practitioner: the diagnostic sequence used to assess a learner’s situation, the prioritization rules that determine which interventions to apply first, the case-pattern recognition that distinguishes expert from novice problem-solving in the domain. This layer corresponds to the externalization of tacit knowledge in the framework of Nonaka and Takeuchi (1995). It is not, by default, present in the training data of general-purpose large language models, because it has typically not been externalized in publicly accessible form.

Layer 3 — Implementation and accountability. The structures through which a learner translates information and methodology into behavior change: implementation intentions (Gollwitzer, 1999), deliberate practice with feedback (Ericsson et al., 1993), retrieval practice (Roediger & Karpicke, 2006), and accountability mechanisms with memory of prior commitments. This layer is, in the language of Pfeffer and Sutton (2000), what closes the knowing-doing gap. It is not delivered by a static video library, and it is not delivered by a general-purpose conversational AI that has no persistent record of the learner.

The *Information–Implementation Divide* is the empirical and economic claim that, in the post-ChatGPT market, the value of a paid online education product is determined increasingly by its position relative to Layers 2 and 3 rather than its position relative to Layer 1. Courses positioned exclusively at Layer 1 face commoditization pressure. Courses that meaningfully integrate Layers 2 and 3 face substantially less commoditization pressure, because the substrate of those layers — externalized proprietary methodology and implementation infrastructure — is not directly substitutable by a general-purpose model.

3.3 Custom GPTs vs. Methodology-Encoded Systems: A Technical Distinction

A natural objection to the framework above is that creators can readily build a “Custom GPT” or comparable lightly customized assistant on top of a general-purpose model, importing some of their own materials through a system prompt or attached documents. If such tools provide a low-cost path to Layer 2, the framework’s predictions about insulation from commoditization may be overstated.

The technical literature on retrieval-augmented generation suggests that this objection underestimates the depth of the divide. Bechard and Marquez Ayala (2024), in research conducted at ServiceNow, documented that generic large-language-model outputs hallucinate at rates incompatible with enterprise or educational use, and that only systems with explicit retrieval grounding meaningfully reduce hallucination. Even with retrieval grounding in place, more recent benchmarking work has shown that state-of-the-art systems struggle to reach 80% factual accuracy on the RAGTruth benchmark without architectures designed explicitly to enforce grounding (Wood et al., 2024). A lightly customized assistant — one that augments a general-purpose model with a small set of documents through a system prompt — typically blends the creator’s externalized methodology with the model’s general training data in ways that the user cannot easily distinguish, and that the creator typically cannot audit. The output is a hybrid: partly the creator’s method, partly the internet’s consensus, and partly the model’s stylistic priors, with no reliable mechanism to enforce which layer dominates a given response.

A *methodology-encoded system*, in the sense developed in this paper, is a system designed from the out-

set to ground its responses in an externalized representation of the creator’s specific methodology, with retrieval, prompting, and evaluation architectures aligned with that constraint. Such systems remain imperfect — no current architecture eliminates hallucination — but the relevant comparison is not “perfect” versus “imperfect” but rather “grounded in the creator’s externalized method” versus “grounded in the open web.” The labor-economic evidence reviewed in §2.2 suggests that this distinction is consequential: Brynjolfsson and colleagues (2025) found the largest productivity gains from AI assistance precisely in conditions where the AI was structured to disseminate the patterns of the most productive workers in the relevant context, not where workers had unstructured access to a generic model.

3.4 Proposed Framework: The *CursoVivo* Implementation Model

The *CursoVivo* implementation model is offered as a concrete instantiation of the *Information–Implementation Divide* framework, designed for the specific case of an existing online course in the Spanish-language creator-economy market. The model embeds a methodology-encoded artificial-intelligence layer within the structure of an existing course, rather than positioning AI as a stand-alone product or a replacement for the course itself. The model is built around six functional components, each mapped to a specific finding from the literature reviewed in §2:

1. **Personalized weekly plans.** Each learner receives an individualized weekly plan, generated by the system from the creator’s externalized methodology and the learner’s current state, specifying objectives, concrete tasks, deadlines, and contingency steps. This component operationalizes implementation intentions (Gollwitzer, 1999) at scale and addresses Bloom’s (1984) personalization challenge through methodology-grounded individualization.
2. **Memory-bearing check-ins.** The system conducts structured weekly check-ins that retain and reference the learner’s prior commitments, completed tasks, and specific obstacles, adjusting subsequent plans accordingly. This component addresses the accountability gap identified by Pfeffer and Sutton (2000) and provides the persistent state that distinguishes a methodology-encoded system from a stateless general-purpose conversational interface.
3. **Deliverable production.** Beyond responding to questions, the system supports the production of concrete deliverables — drafts, checklists, schedules, scripts, routines — adapted to the type of transformation the course is designed to produce. This component operationalizes the principle that learning consolidates through action and external artifacts rather than through passive consumption.
4. **Progress dashboard.** Each learner has access to a structured view of current module, completed work, outstanding commitments, and pending tasks. This component supports the metacognitive monitoring associated with successful self-regulated learning and provides the creator with population-level visibility into completion patterns that information-only courses typically lack.
5. **Daily-action focus.** Each day, the learner is presented with one specific, concrete task, with structured support if the task proves difficult. This component reflects the insight from cognitive-load research that learners benefit from low-friction sequencing of immediate actions rather than from being presented with the full complexity of a multi-week curriculum at once.

- 6. Methodology-grounded conversational layer.** A conversational interface, grounded in the creator's externalized methodology rather than in general web data, that has access to the learner's current week, recent commitments, and completed work. This component addresses the technical distinction developed in §3.3 and operationalizes the AI-as-amplifier-of-expert-method pattern documented by Brynjolfsson and colleagues (2025).

The CursoVivo implementation model is not advanced as a general theory of AI in education. It is advanced as a specific design pattern for course creators in the Spanish-language market who already possess externalizable methodology and an existing course, and who are positioned to benefit from the asymmetric pattern documented in §3.1.

3.5 Practical Implications for Course Creators

The framework developed above yields several specific implications for course creators considering their post-ChatGPT competitive positioning. Three are particularly consequential.

First, the design objective should shift from *information delivery* to *transformation per learner*. The economic logic of the *Information–Implementation Divide* implies that the durable price-defensible value in the post-LLM market is the produced outcome — completed deliverables, sustained behavior change, demonstrable application — rather than the volume of content delivered. The historical practice of measuring course success by content quantity, production value, or initial conversion rate is increasingly disconnected from the structural sources of value in the post-ChatGPT market.

Second, the strategic priority should shift from *acquisition* to *completion*. Reich and Ruipérez-Valiente's (2019) finding that MOOC completion rates have not improved over six years despite continuous platform investment, combined with practitioner data showing that high-implementation cohort programs achieve completion rates approaching one order of magnitude above MOOC averages (Habif, 2017), implies that the largest available efficiency gains for most creators are not in the marketing funnel above the sale, but in the implementation layer below it. A course with a 5% completion rate generates 5% of the testimonials, case studies, and demonstrable outcomes that are themselves the most credible inputs to the marketing funnel; doubling completion is, in this view, equivalent to doubling marketing efficiency, with potentially favorable economics.

Third, defensive responses that compete directly with ChatGPT on the information dimension are unlikely to succeed. Reducing prices to compete with a free, near-infinite information source is not a sustainable strategy in any market governed by the economics of information goods (Shapiro & Varian, 1999). The strategic alternative — investing in the externalization of proprietary methodology and the construction of methodology-grounded implementation infrastructure — is a non-trivial undertaking, but it is positioned in the layers of the market that the structural analysis in §3.2 identifies as defensible.

4. Conclusions

4.1 Summary of Findings

This paper has examined whether the prevailing narrative that ChatGPT is killing the online course market is supported by the available empirical literature. The evidence reviewed suggests that the narrative is, at best, incomplete. The completion crisis in self-paced online education predates ChatGPT by more than a decade and reflects a structural feature of information-only delivery (Reich & Ruipérez-Valiente, 2019; Bloom, 1984). The commoditization of explicit information was already predicted by the economics of information goods more than two decades before generative AI reached consumer scale (Shapiro & Varian, 1999) and has been accelerated, but not invented, by recent reductions in the marginal cost of model inference (Stanford HAI, 2024). Behavior-change research consistently demonstrates that information transfer alone does not produce skill acquisition or sustained action; implementation intentions (Gollwitzer, 1999), deliberate practice (Ericsson et al., 1993), retrieval practice (Roediger & Karpicke, 2006), and accountability structures (Pfeffer & Sutton, 2000) are the mechanisms that produce outcomes, and none of them are delivered by a static video course or by a general-purpose conversational AI.

The aggregate market evidence is consistent with this structural analysis. Information-services firms positioned at the commodity-information layer have contracted; creator-economy platforms positioned at the implementation and community layer have continued to grow during the same period (Hotmart Company, 2024; Kajabi, 2025; Sheridan, 2023). The asymmetric pattern is what the economics of information goods and the labor economics of generative AI jointly predict.

ChatGPT, in this reading, is not a discontinuity in the online education market. It is a clarifying event that has made visible a structural divide that was already present. The *Information–Implementation Divide* — the proposition that the durable value of paid online education in the post-LLM market is determined by position relative to the implementation and methodology layers rather than the information layer — is offered as an integrative framework for further research and practical design. The CursoVivo implementation model is offered as one concrete instantiation of the framework. The competitive frontier in adult digital education is not access to information; it is the production of measurable implementation grounded in proprietary methodology.

4.2 Limitations

This review is subject to the selection bias inherent in narrative literature reviews. The framework proposed has not yet been validated through controlled experimental studies, and the practitioner-reported completion data cited from cohort-based programs (Habif, 2017) reflect platform self-reporting rather than independent measurement. Several widely circulated industry statistics regarding specific platform performance and market growth rates could not be traced to primary sources and were excluded from the analysis; the resulting claims are therefore more conservative than those that appear in the popular literature. The analysis is concentrated on the adult digital-education market and may not generalize to formal accredited education, K–12 contexts, or workforce-training programs governed by regulatory frameworks. Finally, the technical literature on retrieval-augmented generation and methodology-grounded AI systems is evolving rapidly; specific architectural claims may require revision as the field develops.

4.3 Future Research Directions

Three directions for subsequent research follow naturally from the framework. First, controlled comparative studies of completion rates and learning outcomes between matched cohorts using information-only courses, lightly customized general-purpose AI tools, and methodology-encoded implementation systems would provide direct empirical tests of the *Information–Implementation Divide*. Second, longitudinal studies of creator-economy platform performance — disaggregated by the position of platform offerings within the three-layer framework developed in §3.2 — would clarify the rate at which the asymmetric pattern is consolidating. Third, the externalization of tacit methodology into machine-readable form, building on the SECI framework of Nonaka and Takeuchi (1995), is a substantial methodological challenge in its own right and represents a productive intersection of knowledge management, instructional design, and applied AI research. Empirical research grounded in the Latin American creator economy — a globally significant and academically underrepresented segment of the digital education market — would be a particularly valuable extension.

Resumen en Español

El relato predominante en la economía de creadores sostiene que ChatGPT y los modelos de lenguaje de gran escala están acabando con el mercado de cursos online. Casos visibles como el colapso bursátil de Chegg y encuestas que muestran a estudiantes sustituyendo cursos pagos por ChatGPT han reforzado esta visión. Este paper sostiene que el relato es empíricamente incompleto y estructuralmente engañoso. Apoyado en investigación revisada por pares en educación, psicología del comportamiento, gestión del conocimiento y economía laboral de la IA generativa, el paper documenta que (a) la crisis de completación en cursos auto-administrados es muy anterior a ChatGPT, (b) la commoditización de la información explícita ya estaba predicha por la economía de bienes de información dos décadas atrás, y (c) la frontera competitiva post-ChatGPT en educación adulta está en la *brecha de implementación* — el hecho documentado de que la sola transferencia de información no produce cambio de conducta. La evidencia de mercado en plataformas de creadores cuya propuesta de valor está centrada en implementación, comunidad y metodología — y no en entrega de información — respalda esta tesis: estas plataformas continúan creciendo durante el mismo periodo en el que servicios solo-información se han contraído. El paper propone la *Brecha Información–Implementación* como marco integrador e introduce el modelo CursoVivo — una capa de IA entrenada específicamente en la metodología del creador — como instanciación concreta. Se discuten implicaciones para creadores, investigadores EdTech y diseño de sistemas de aprendizaje digital post-LLM.

Palabras clave: ChatGPT y cursos online; impacto de la IA en educación; disrupción del e-learning; estrategia para creadores de cursos; información como commodity; brecha de implementación; IA en educación online; CursoVivo; tasas de completación; conocimiento tácito; MOOC; IA generativa en educación; diseño de aprendizaje post-LLM.

References

- Acuña Caicedo, R. M., Gómez Soto, J. M., López Batista, V. F., & Mendoza-Moreno, M. A. (2023). Revolutionizing online learning: The potential of ChatGPT in massive open online courses. *European Journal of Education and Pedagogy*, 4(6). <https://eu-opensci.org/index.php/ejedu/article/view/30686>
- Bechard, P., & Marquez Ayala, O. (2024). Reducing hallucination in structured outputs via retrieval-augmented generation [Preprint]. *arXiv*. <https://arxiv.org/abs/2404.08189>
- 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>
- 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>
- 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>
- Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 493–503. <https://doi.org/10.1037/0003-066X.54.7.493>
- Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119. [https://doi.org/10.1016/S0065-2601\(06\)38002-1](https://doi.org/10.1016/S0065-2601(06)38002-1)
- Habif, S. (2017, December 21). How to design an online course with a 96% completion rate. *Psychology of Stuff*. <https://medium.com/behavior-design/how-to-design-an-online-course-with-a-96-completion-rate-180678117a85>
- Hosseini, M., Gao, C. A., Liebovitz, D. M., Carvalho, A. M., Ahmad, F. S., Luo, Y., MacDonald, N., Holmes, K. L., & Kho, A. (2023). An exploratory survey about using ChatGPT in education, healthcare, and research. *PLOS ONE*, 18(10), e0292216. <https://doi.org/10.1371/journal.pone.0292216>
- Hotmart Company. (2024, March 12). *Hotmart Company, home to Teachable, announces record-breaking \$10 billion in global creator earnings* [Press release]. <https://press.hotmart.com/hotmart-company-announces-record-breaking-10-billion-in-global-creator-earnings>
- 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/>
- Kajabi. (2025, August 6). *\$10B in creator revenue and climbing: Kajabi creators are achieving long-term financial success* [Press release]. Business Wire. <https://www.businesswire.com/news/home/20250806424975/en/>
- McKinsey & Company. (2024). *The state of AI in early 2024: Gen AI adoption spikes and starts to generate value*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>
- Mollick, E., Mollick, L., Bach, N., Ciccarelli, L. J., Przystanski, B., & Ravipinto, D. (2024). AI agents

and education: Simulated practice at scale [Preprint]. *arXiv*. <https://arxiv.org/abs/2407.12796>

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.

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

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

Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>

Shapiro, C., & Varian, H. R. (1999). *Information rules: A strategic guide to the network economy*. Harvard Business School Press.

Sheridan, E. (2023, April 19). *The creator economy could approach half-a-trillion dollars by 2027*. Goldman Sachs. <https://www.goldmansachs.com/insights/articles/the-creator-economy-could-approach-half-a-trillion-dollars-by-2027>

Sherman, A. (2023, May 2). Chegg shares drop more than 40% after company says ChatGPT is killing its business. *CNBC*. <https://www.cnbc.com/2023/05/02/chegg-drops-more-than-40percent-after-saying-chatgpt-is-killing-its-business.html>

Stanford Institute for Human-Centered AI. (2024). *The 2024 AI Index report*. Stanford University. <https://hai.stanford.edu/ai-index/2024-ai-index-report>

Wood, M. J., et al. (2024). 100% elimination of hallucinations on RAGTruth for GPT-4 and GPT-3.5 Turbo [Preprint]. *arXiv*. <https://arxiv.org/abs/2412.05223>