Entrepreneur & Innovation Exchange
Published online at EIX.org on November 19 2026
DOI: 10.32617/1459-6afeeb32599de

Is AI Making Entrepreneurs Overconfident?

Ernest Benchek III (University of Pittsburgh School of Business)

KEYWORDS: Artificial Intelligence, Entrepreneurship.

AI can help entrepreneurs create impressive business plans without understanding their markets. Here's how to distinguish real readiness from false confidence.

A student or aspiring entrepreneur can now walk into an entrepreneurship classroom, accelerator, incubator, or pitch competition carrying a polished business plan, market analysis, customer persona, financial projection, launch strategy, and investor deck, all generated in an afternoon. The work looks sophisticated. The slides are clean, the assumptions plausible, the financial model professional. Yet the founder may not understand the customer, the market, the operational constraints, or the risks buried in the proposal.

This is a new problem for entrepreneurship practice. Generative AI can make founders look prepared, and, more dangerously, feel prepared, before they are. What follows separates *artifact fluency*, the ability to produce a venture's plans and decks, from *founder readiness*, the ability to build the venture.

The problem is not that founders use AI. They will, and often they should. AI can help them brainstorm, test assumptions, compare business models, structure customer discovery, analyze competitors, and communicate more clearly. Used well, it lowers barriers to entrepreneurship and helps capable people move from hesitation to action.

The danger sits on the other side of that same capability. Generative AI may raise *perceived* entrepreneurial capability faster than it raises *actual* entrepreneurial capability. It can make people feel more ready to start a venture without making them more able to build one. That gap matters, because entrepreneurship is not entered with words on a page. It is entered with time, money, relationships, reputation, opportunity cost, and real execution risk.

AI and self-delusion

The problem is fundamentally one of entrepreneurial calibration: knowing which beliefs about one's own capabilities and about the venture rest on understanding and testing, and which AI has merely made look plausible. A calibrated founder holds confidence in proportion to what they understand, can demonstrate, and have reason to believe, without the overconfidence that AI can bring to founders who mistake its output for their own understanding. That is because AI can expand what founders are able to accomplish without necessarily expanding what they personally understand. An accurate financial model, for example, can make financial competence seem greater than it is. This is not a criticism of relying on tools or outside expertise. The danger arises when founders mistake the capabilities they access for capabilities they personally possess.

AI can also inflate perceived product-market fit. A polished market analysis, customer profile, or revenue projection may make demand appear established without adding evidence that customers want the product. AI can help interpret market information, but it cannot replace direct customer engagement or actual customer and operational data. Founders must distinguish a plausible account of demand from evidence of demand.

The gap is not accidental

This is not a hypothetical worry. It sits at the intersection of several well-established research traditions (on entrepreneurial confidence, on overconfidence, and on how people use automation) and a fast-growing literature on generative AI in entrepreneurship.

The first is entrepreneurial self-efficacy: a person's belief that they can perform the tasks entrepreneurship requires. Building on Bandura's foundational work on self-efficacy, Chen, Greene, and Crick (1998) helped establish entrepreneurial self-efficacy as a distinct entrepreneurial belief, finding that business founders reported higher self-efficacy in innovation and risk-

Copyright © 2026 The Authors. Entrepreneur & Innovation Exchange is published at EIX.org. This is an open access article under the terms of the Creative Commons Attribution-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited and no modifications or adaptations are made. View EIX.org Authorship Terms at https://eix.org/terms

taking than nonfounders. Subsequent reviews link entrepreneurial self-efficacy to entrepreneurial intentions, behavior, and other outcomes (Newman et al., 2019). But self-efficacy is a belief about capability, and beliefs can be miscalibrated. In their hubris theory of entrepreneurship, Hayward, Shepherd, and Griffin (2006) argue that highly confident actors may be moved to start ventures despite high failure rates and then allocate resources based on confidence that exceeds the available evidence.

A second body of research explains how AI can widen that gap. Human-factors work shows that people can over-trust automated and decision-support systems, monitor them less closely, and accept their recommendations with insufficient verification, patterns Parasuraman and Manzey (2010) describe as automation complacency and automation bias. Generative AI is the most fluent automation ever placed in front of aspiring founders. It produces exactly the artifacts entrepreneurship has always used to signal progress, and it produces them well enough to lower the founder's own scrutiny.

Recent research extends this logic directly to entrepreneurship. Studies suggest generative AI can raise entrepreneurial self-efficacy and intention (Xie & Wang, 2025), while integrative and conceptual work describes a double edge: the same tools that empower founders by lowering barriers may also entrap them by inflating overconfidence, weakening critical evaluation, or making venture ideas appear more plausible than their evidence warrants (Lu, Zhao, & Zheng, 2026; Madjdi & Wurth, 2026).

Put these together and the risk is structural, not incidental. AI raises the input to entrepreneurial confidence (polished, capable-looking output) while quietly suppressing the scrutiny that would normally keep that confidence honest. The founder ends up feeling more ready precisely because the tool did the part they could not yet do themselves.

A polished pitch does not signal founder readiness

Entrepreneurship programs have long judged founder progress through artifacts: business plans, pitch decks, customer profiles, competitive analyses, revenue models, financial projections. Those artifacts still matter. But AI has changed what they signal.

A founder who produces a polished pitch deck may have strong entrepreneurial judgment, or may simply have written a good prompt. A founder who presents a detailed market analysis may understand the market deeply, or may have accepted an AI summary without knowing whether it reflects real customer behavior. A founder who shows five-year projections may grasp the unit economics, or may have numbers that look credible and collapse under one round of questioning.

The artifact has come loose from the capability it used to imply. Educators, accelerators, and investors can no longer treat artifact quality as a reliable proxy for founder readiness. The two need to be named and separated.

- *Artifact fluency* is the ability to produce entrepreneurship-like outputs: plans, decks, analyses, personas, projections, narratives. Generative AI is very good at helping people produce these.
- *Founder readiness* is different. It is the ability to create opportunities by validating demand, understanding customers, making tradeoffs under uncertainty, managing cash, selling, absorbing feedback, assembling resources, executing under pressure, and knowing when to pivot or stop.

AI can support both. But it can raise the appearance of the first without meaningfully building the second. That is the core risk this article is about.

Distinguishing artifact fluency from founder readiness therefore requires examining how founders reach and assess their conclusions, not simply what they produce. Four questions can help reveal whether AI has strengthened judgment or merely increased confidence: Can founders explain their assumptions, show what they have learned from real people, identify what AI actually improved, and say what would change their minds?

AI's upside

The story is not one-directional, and it would be a mistake to treat AI only as a confidence inflator. The same tool can sharpen judgment. Used skeptically, AI can expose weak assumptions, unrealistic projections, poor market fit, or thin customer evidence before a founder commits. It can talk a founder out of a venture that will not survive, and doing so early and cheaply is a

Copyright © 2026 Ernest Benchek III, Published by Entrepreneur & Innovation Exchange

DOI: 10.32617/1459-6afeeb32599de

service, not a failure.

So the useful question is not whether AI produces more entrepreneurship or less. It is whether AI improves *calibration*, whether it helps the right founders enter for the right reasons, and helps the wrong ventures surface their flaws before real resources are at stake.

A framework: AI's four effects on venture entry

AI can shape the decision to enter a venture along two dimensions: whether the founder and venture are genuinely ready, and whether AI pushes the founder toward entry or away from it. Crossing them yields four outcomes.

	AI encourages entry	AI discourages entry
Founder / venture is ready	**Calibrated entry**	**Missed opportunity**
Founder / venture is not ready	**Premature entry**	**Productive deterrence**

Calibrated entry is the best case. A capable but hesitant founder, someone who already understands the problem, customer, or technology, uses AI to organize their thinking, communicate the opportunity, and move faster. AI removes friction and enables productive action that might otherwise have stalled.

Productive deterrence is the quiet win. AI helps a founder see that a venture is not ready. The problem is vague, the market too small, the margins fantasy, the competitive position weak, or the resources absent. Discouraging that entry is a good outcome, delivered before the founder has spent a year finding out the hard way.

Premature entry is the failure mode this article exists to name. An underprepared founder produces convincing artifacts, receives encouraging feedback on them, and begins committing time and money before validating the venture or building the capacity to execute. This is where automation complacency meets entrepreneurial hubris, and it is the single most important risk for educators, accelerators, and early-stage investors to manage.

Missed opportunity **and** ***over-deterrence*** are the mirror risks on the discouragement side, and they are worth telling apart. Missed opportunity is about the *founder*: A genuinely capable person is talked out of acting by an over-cautious AI reading. Over-deterrence is about the *venture*: A promising but unconventional idea is judged weak because it does not match historical patterns, and some good ventures look bad precisely until the founder creates the market, educates customers, or builds legitimacy that no training data yet reflects. AI trained on the past will systematically underrate the genuinely new.

The goal is not more entrepreneurship or less. It is better-calibrated entrepreneurship: AI that encourages prepared founders, disciplines underprepared ones, and helps everyone test assumptions rather than merely dress them up.

New rules for educators, accelerators, and investors

If a polished artifact no longer signals readiness, the people who evaluate founders need a different question. The old one was: *Is the business plan good?* The new one is: *Does the founder understand and own the assumptions behind it?*

1. Can the founder explain and defend the assumptions behind the AI output?

If a founder cannot explain their market size, pricing logic, customer problem, margins, or go-to-market plan without simply repeating the AI output, they may not own the judgment behind it. A useful exercise is to close the deck and ask them to reconstruct the decision: How did you arrive at that market size? Which customers are included, and why? What evidence supports the price? Which assumptions are most uncertain, and which could kill the business? The founder need not recreate every technical task that AI performed, but should understand enough to judge its implications and limitations.

The point is not to catch AI use. It is to find unowned judgment. A founder who used AI well will answer these better, not worse. The tool should have deepened their understanding, not substituted for it. If the answers exist only inside the slides, that is the tell.

2. Has the founder tested the venture against real people?

AI-generated customer personas are not customers. AI-generated pain points are not customer discovery. AI-generated market summaries are not demand validation. This is where the lean-startup tradition remains relevant in an AI world: Blank's (2013)

Copyright © 2026 Ernest Benchek III, Published by Entrepreneur & Innovation Exchange

DOI: 10.32617/1459-6afeeb32599de

approach emphasizes testing business assumptions through direct customer feedback rather than relying on static business plans. Founders should be able to say how many customers, users, buyers, partners, or industry participants they have actually spoken with; what surprised them; which original assumptions turned out to be wrong; and what those conversations told them that their analysis, AI-assisted or not, did not anticipate.

The standard is the same whether or not AI was involved: Has the venture met reality yet? A founder who has used AI well arrives at those conversations with sharper questions, not with an excuse to skip them. And if stakeholder conversations have not changed the venture in any meaningful way, that is a warning sign in itself, a signal that the founder may be using AI, or any planning tool, to confirm the idea rather than test it.

3. Did AI improve the venture, or only the presentation?

A better deck is not a better business. A cleaner model is not validated demand. More professional projections are not financial discipline. The diagnostic question is simple: *What actually changed because of AI?* Did the founder discover a better customer segment, identify a real weakness, revise the pricing, abandon a bad assumption, narrow the market, or cut the capital required? Or did the tool simply make the same venture sound more convincing? If AI improved only persuasiveness, it likely raised confidence without raising readiness, the exact premature-entry signature.

4. What evidence would make the founder pivot or stop?

Prepared founders know what would change their mind. Overconfident ones keep going because the story still sounds good — the hubris pattern in action. Every plan should carry explicit decision criteria: what evidence justifies continuing, what would force a pivot, and what would mean stopping before more is spent. But the criteria have to be defensible. “If we don't reach 3% penetration in 18 months, we shut down” is only useful if the founder can say why 3%, why 18 months, and why missing it means the opportunity is gone. Is the threshold tied to runway, customer acquisition cost, adoption rates, investor milestones, or competitive dynamics? AI should help founders build and stress-test these tripwires — not manufacture more convincing reasons to keep going.

How founders can test their judgment

The solution is not to discourage founders from using AI. It is to help them distinguish what AI enables them to produce from what they actually understand. Founders can do this through three practices that focus on AI-assisted judgment, rather than repeating the venture-validation steps every entrepreneur should take.

1. **Identify where AI is substituting for understanding**. After generating a plan, model, or analysis, founders should distinguish conclusions they can explain and defend from those they accept mainly because AI produced them. Which assumptions do they understand? Which judgments depend on expertise they do not yet possess? The question is not whether AI did the work, but whether the founder recognizes the boundary between the tool’s contribution and their own competence.
2. **Challenge AI-generated conclusions rather than simply accepting them**. Ask the tool to expose weak assumptions, alternative interpretations, and consequences that the initial answer overlooked. Instead of “make this idea sound better,” the prompts that build readiness are the adversarial ones: *What would make this venture fail? Which assumptions are weakest? What evidence should I gather before spending money? What would a sharp investor attack? What customer behavior would disprove this? What would tell me to pivot, or to stop?* Then assess those critiques independently; a second fluent AI answer is not, by itself, independent confirmation. The goal is to make AI a discipline mechanism, not just a confidence engine.
3. **Recalibrate confidence based on actual understanding.** Did working with AI clarify tradeoffs and improve the founder’s ability to decide, or did it mainly produce a more persuasive artifact? A founder who recognizes a gap can learn the relevant skill, enlist expertise, or explicitly account for dependence on the tool. Entrepreneurship has always involved assembling resources. Good calibration means knowing what the founder can judge, what AI can contribute, and where neither should be trusted without further evidence.

Copyright © 2026 Ernest Benchek III, Published by Entrepreneur & Innovation Exchange

DOI: 10.32617/1459-6afeeb32599de

Takeaways

Evaluators have matching practices. Educators can require an *assumption audit* alongside any AI-assisted work, a short document identifying which claims came from AI, which were externally validated, which remain speculative, and which are most dangerous if wrong. Accelerators can run *AI defense interviews*, in which founders explain and defend the plan with the slides closed, separating presentation quality from understanding. Investors can ask, pointedly, "What did AI get wrong?" A founder who used the tool critically can always name where it was too optimistic, too generic, or too detached from the market. Venture competitions can rewrite their rubrics to reward evidence of customer contact, assumption testing, financial realism, and demonstrated learning rather than the most polished deck in the room.

The artifacts are now cheap. The capabilities are not. That shift redraws the job. Educators can ask students to explain the judgments behind AI-assisted work; accelerators can probe which decisions founders understand and which depend on the tool; investors can examine the reasoning behind the deck rather than its polish. For founders, the task is to identify where AI substitutes for understanding, challenge its conclusions, and recalibrate confidence accordingly. These practices complement, rather than replace, ordinary customer discovery and venture validation.

AI can widen access to entrepreneurship, and that is worth wanting. It can expand what founders accomplish without necessarily expanding what they personally understand. That distinction does not diminish AI's value: Entrepreneurs have always drawn on capabilities beyond their own. But confidence becomes dangerous when founders mistake borrowed capability for personal competence, or plausible analysis for justified judgment. Used well, AI can help founders recognize those boundaries, challenge weak assumptions, and move with justified speed. The promise of AI-enabled entrepreneurship is not more founders, faster pitches, or better-looking plans. It is better-calibrated entrepreneurial action.

Which leaves every educator, accelerator, investor, and founder with one question to ask of any AI-assisted venture:

Did AI make this founder more prepared, or only make them feel that way?

References

Bandura, A. (1997). Self-Efficacy: The Exercise of Control. New York: W. H. Freeman.

Blank, S. (2013). Why the lean start-up changes everything. Harvard Business Review, 91(5), 63–72.

Chen, C. C., Greene, P. G., & Crick, A. (1998). Does entrepreneurial self-efficacy distinguish entrepreneurs from managers? Journal of Business Venturing, 13(4), 295–316.

Hayward, M. L. A., Shepherd, D. A., & Griffin, D. (2006). A hubris theory of entrepreneurship. Management Science, 52(2), 160–172.

Lu, J. G., Zhao, G. G., & Zheng, A. M. (2026). Generative AI use in entrepreneurship: An integrative review and the Empowerment–Entrapment Framework. Strategic Entrepreneurship Journal. Advance online publication. https://doi.org/10.1002/sej.70050

Madjdi, F., & Wurth, B. (2026). AI-mediated plausibility regimes: Entrepreneurial judgment, epistemic risk, and the distribution of entrepreneurial futures. Journal of Business Venturing Insights, 26, Article e00644.

Newman, A., Obschonka, M., Schwarz, S., Cohen, M., & Nielsen, I. (2019). Entrepreneurial self-efficacy: A systematic review of the literature on its theoretical foundations, measurement, antecedents, and outcomes, and an agenda for future research. Journal of Vocational Behavior, 110, 403–419.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381–410.

Xie, Y., & Wang, S. (2025). Generative artificial intelligence in entrepreneurship education enhances entrepreneurial intention through self-efficacy and university support. Scientific Reports, 15, Article 24079.

Copyright © 2026 Ernest Benchek III, Published by Entrepreneur & Innovation Exchange

DOI: 10.32617/1459-6afeeb32599de