

EOS RESEARCH REPORT

Built Before It Is Real

AI, abundance, and the widening gap
between production and market proof





About this report

This report examines a change that is easy to see but harder to interpret. Generative AI has sharply reduced the time, skill, and cost required to produce many of the visible components of a business. A founder can now move from an idea to working copy, code, analysis, visual assets, sales material, customer-service workflows, research summaries, and prototypes with far fewer people and far less elapsed time than was previously possible. The resulting gain in access is real, and much of the strongest available evidence confirms significant improvements in speed and productivity.

The question explored here is what happens after that gain. The report argues that the production side of entrepreneurship is becoming abundant faster than the market side. The ability to make something credible has become easier to acquire, while the work of finding demand, earning trust, learning from customers, reaching buyers repeatedly, building distribution, and turning early interest into durable revenue remains difficult. The argument is therefore not that AI has made entrepreneurship worse, or that AI-assisted firms fail more often. The available evidence does not support those claims. The more defensible proposition is that AI is changing the **relative scarcity** of the capabilities required to build a successful company.

The report draws on peer-reviewed research and official statistical sources. It avoids treating working papers, vendor surveys, arXiv preprints, or unsupported commercial claims as established evidence. Several concepts introduced here — including the **Production–Market Divide**, **Entrepreneurial Scarcity Shift**, **Closed-Loop Founder**, **Market Evidence Ladder**, and **Market-Embedded AI Company** — are EOS analytical frameworks developed from the evidence reviewed. They are intended to help founders and operators think more clearly about where AI creates leverage and where it can create false confidence; they are not presented as validated predictive models.

Enterprise Operational Systems | September 2026 | eospk.com

Contents

Executive Summary	3
Introduction	6
Chapter 1	7
Chapter 2	10
Chapter 3	12
Chapter 4	15
Chapter 5	18
Chapter 6	20
Chapter 7	23
Chapter 8	26
Chapter 9	28
Chapter 10	30
Chapter 11	32
Diagnostic	33
Chapter 12	35
Evidence boundaries	37
References	38
Official data sources	39

Page numbers exclude the cover.

Executive Summary

Generative AI has changed the practical threshold for starting a business. Tasks that once required several specialists, outside agencies, or long periods of self-teaching can now be attempted by one person using a relatively inexpensive set of tools. A founder can develop positioning, write a website, produce a pitch, create graphics, research competitors, build a prototype, draft a financial model, automate administrative work, and prepare weeks of marketing material from the same laptop. The important shift is not that every output is exceptional. It is that the cost of reaching a usable first version has fallen enough to change what one person or a very small team can attempt.

The productivity evidence is substantial. Noy and Zhang's controlled experiment found that ChatGPT reduced the time required for professional writing tasks by **40%** while raising average output quality by **18%** (Noy & Zhang, 2023). Cui and colleagues, working across three randomized field experiments involving **4,867** developers, found that access to AI coding assistance increased completed tasks by **26.08%** (Cui et al., 2026). Brynjolfsson, Li, and Raymond studied **5,172** customer-support agents and found that generative AI increased successfully resolved issues per hour by about **15%** on average, with considerably larger gains among less experienced workers (Brynjolfsson, Li, & Raymond, 2025). In each setting, AI reduced the amount of time or expertise required to produce competent work.

Business adoption is also broad enough that these effects can no longer be treated as experimental curiosities. The Federal Reserve's 2026 Small Business Credit Survey found that **46%** of surveyed U.S. small employer firms were already using AI and another **15%** intended to use it within twelve months. Among firms already using AI, **83%** reported using it for writing or marketing, **61%** for individual productivity, and **51%** for planning or analysis; **71%** reported higher productivity, **39%** improved quality, and **31%** increased sales. Official data from the European Union, Canada, and the United Kingdom show the same broad direction even though the surveys use different definitions and populations.

Those figures describe a rapid expansion of production capability. They do not describe a comparable disappearance of the commercial problems that follow production. In the same Federal Reserve survey, **57%** of small employer firms named reaching customers and growing sales as an operating challenge, making it the most common challenge reported. Forty-eight percent cited weak sales as a financial challenge. Among early-stage nonemployer firms that intended to hire, **66%** reported difficulty reaching customers or growing sales. These statistics do not show that AI caused the difficulty; customer acquisition has always been hard. They do, however, show that better production tools can coexist with persistent weakness on the demand side of the business.

This distinction is the starting point for the report. A company requires at least two broad capabilities. The first is the capacity to produce: to create the product, the brand, the copy, the software, the research, the sales material, and the operating systems that allow a company to function. The second is the capacity to become established in a market: to find customers, understand how they make decisions, build trust, learn from objections, create distribution, convert interest into transactions, retain buyers, and develop a reputation. Generative AI is reducing the cost of the first capability very quickly. There is much less evidence that it has reduced the second at the same rate.

The implication is not that production has become unimportant. It is that production may become less distinctive when similar capability becomes widely available. Doshi and Hauser's experiment on generative AI and creative writing illustrates the mechanism. Participants using AI produced stories that were rated as more novel at the individual level, yet the pool of AI-assisted stories became more similar overall (Doshi & Hauser, 2024). The study was not about startups, so it should

not be extended beyond what it measured, but the pattern is relevant: individual quality can rise while the field becomes more homogeneous. In commercial settings, a similar dynamic would mean that better copy, cleaner design, stronger research summaries, and credible-looking products become easier for many competitors to produce at once.

That creates what this report calls the **Entrepreneurial Scarcity Shift**. When production capability is expensive, the ability to produce is itself a meaningful barrier to entry. When production becomes cheaper and more widely distributed, other scarce resources become relatively more important. These include repeated access to the right customers, proprietary market knowledge, reputation, trust, relationships, distribution, high-quality feedback, and the ability to distinguish genuine demand from weak signals of interest. A capability does not need to become harder in absolute terms to become more important; it can become relatively more valuable because another capability has become abundant.

The structural difficulty of moving from entry to scale reinforces this point. The U.S. Census recorded **578,926** seasonally adjusted business applications in July 2026, but an application for an Employer Identification Number is not the same thing as an operating or successful company. For that application cohort, the Census projected **29,959** employer startups with payroll tax liabilities would form within four quarters. OECD evidence on micro-startups shows an even steeper long-run filter: on average, only about **4%** grow, and in a stylized cohort of 100 micro-startups observed after five years, only one to eight reach ten or more employees. Much of this evidence predates the current generative-AI wave, which is exactly why it should be treated as a baseline rather than as proof of an AI effect. The difficulty of scaling was already present; AI is changing how quickly and cheaply a founder can arrive at the point where that difficulty begins.

The report calls the separation between internal production and external commercial learning the **Production–Market Divide**. Production progress occurs when the company creates something: a site, feature, campaign, workflow, analysis, pitch, or automation. Market progress occurs when the company learns something reliable because a real customer, prospect, partner, channel, or competitor has responded. The two can reinforce each other, but they can also drift apart. A firm can publish more without understanding why prospects do not buy, automate outreach before learning which segment has urgency, or optimize a conversion funnel that is attracting the wrong audience. In each case, internal output improves while the commercial model remains uncertain.

Generative AI can widen this divide because it makes internal work unusually easy to continue. There is always another version of the homepage, another segment description, another content plan, another pricing model, another automated sequence, another analysis of the same problem. This creates an **Optimization Trap** in which the organization remains productive without forcing its assumptions into situations where they can be contradicted. The risk is not laziness; it is productive activity directed at questions that the market has not yet answered.

A company becomes especially vulnerable when its information loop closes in on itself. In a **Closed-Loop Founder** model, founder assumptions are elaborated by AI, converted into output, measured through internal or platform metrics, and then interpreted by AI again. That process can look data-driven while still resting on a weak original understanding of the market. A stronger model is open to external correction: customer behaviour produces observations, those observations change interpretation and decisions, the company builds or adjusts something, and the next transaction attempt generates new evidence. AI can be extremely useful inside this loop, but its role is to accelerate analysis and testing rather than substitute for market contact.

The entrepreneurship literature gives strong support to this emphasis on learning. Randomized trials by Camuffo and colleagues show that founders trained to formulate and test assumptions more scientifically become more willing to terminate weak ideas and search more efficiently

(Camuffo et al., 2020; 2024). Koning, Hasan, and Chatterji found substantial performance gains among startups adopting A/B testing, while Koning and Nanda showed how experimentation can still mislead when the test audience is badly matched to the intended customer. Song, Wang, and Parry found that formal collection and use of market information mattered more for venture performance than customer interaction in the abstract. The common lesson is that contact with the market becomes valuable when it is converted into disciplined learning and allowed to change decisions.

Relationships are part of this learning system because markets are social as well as informational. Stam, Arzlanian, and Elfring's meta-analysis found a positive association between entrepreneurial network social capital and small-firm performance. Palmatier and colleagues' meta-analysis found that relationship quality and relationship investment can materially influence performance, particularly in B2B, services, and channel contexts. Research on adaptive selling similarly shows that sales performance improves when sellers respond to the customer rather than rely on a fixed script. None of this implies that networking is a substitute for a good product. It shows that access, trust, referral, interpretation, and market knowledge often move through relationships that cannot be recreated simply by generating more content.

The authenticity literature adds another layer. Experiments on AI-authored marketing suggest that consumers can penalize AI involvement in contexts where communication is expected to carry a human emotional or personal signal, while the same penalties are weaker when information is factual, AI assists rather than replaces a human, or AI use is already expected. The useful conclusion is not that customers dislike AI. It is that automation should be designed according to what the interaction is meant to convey. A routine status update and an apology after a serious failure are not the same communication problem, even if both can technically be automated.

The report therefore does not argue for less AI. It argues for a different standard of AI use. The relevant question is whether a new tool increases the company's exposure to market reality or mainly increases the volume of internally generated work. AI can summarize interviews, compare sales calls, cluster support tickets, prepare account research, identify repeated objections, analyze churn, generate test variants, and remove administrative tasks that keep experienced people away from customers. Used this way, it shortens the path between market evidence and organizational response.

The final sections of the report develop a practical model for doing this deliberately. The **Market-Embedded AI Company** is evaluated across six dimensions: production velocity, customer-contact velocity, experiment velocity, learning conversion, relationship density and distribution access, and revenue evidence. Together, these dimensions shift attention away from how much the organization can produce and toward how quickly it turns production into reliable market learning. The central claim of the report is therefore narrower, and more useful, than the common argument that AI has made entrepreneurship easier. AI has made it much easier to assemble the supply side of a business. The competitive question now is whether founders use that leverage to become more deeply connected to demand, or simply to create a more polished version of an untested idea.

Introduction — A Company by Sunday Night

It is Friday evening and you have an idea for a company. It may be a software tool, a research service, a specialist agency, a health product, an ecommerce brand, a consultancy, or something less defined. A few years ago, the distance between that idea and a credible launch would have been determined by the people, skills, money, and time you could assemble. Even a small launch required a sequence of tasks that rarely sat inside one person's skill set: market research, positioning, writing, design, development, financial modelling, customer communication, basic legal documentation, analytics, and a growing list of operational details.

By Sunday night, a surprising share of that work can now be attempted by one person. An AI model can help map the market, propose segments, draft positioning, write the website, generate product descriptions, structure a financial model, summarize competitor reviews, create interview questions, prepare sales emails, and assist with a working prototype. Connected to image generators, code assistants, payment tools, analytics, and automation software, the same founder can move from a vague concept to something that has the outward shape of a business in a matter of days.

This is a genuine expansion of access to entrepreneurship. It lowers the cost of experimentation, allows people to attempt tasks that would previously have been deferred or outsourced, and reduces the minimum team required to make a company look operational. The effect is particularly important for people who have a strong understanding of a problem but lack one of the specialist capabilities that used to stand between expertise and execution. A domain expert can now draft software, a technical founder can produce competent sales material, and a small operator can perform basic analysis without hiring a dedicated analyst.

The difficulty begins when the visible completion of the business is mistaken for commercial completion. A website can be excellent without proving that the intended buyer cares. A polished pitch can still target the wrong decision-maker. A customer persona can be coherent while missing the language customers use themselves. A content operation can run efficiently without providing a repeatable route to demand. AI can generate plausible answers to each of these questions, but plausibility and evidence are different things. The market remains the place where the company learns whether its assumptions survive contact with people who are free to ignore it, compare it, reject it, negotiate with it, pay for it, or leave.

The central question of this report follows from that distinction: **what happens to entrepreneurship when the cost of making a business falls much faster than the cost of being discovered, trusted, chosen, paid, and retained?** If production continues to become cheaper while market access remains scarce, the advantage of the next generation of firms may depend less on whether they can build and more on whether they can use faster building to learn from the market faster than their competitors.

CHAPTER 1

The Company You Can Build in a Weekend

The most immediate effect of generative AI is often described as productivity, but that word understates what has changed for a small firm. Traditional productivity tools made an existing task more efficient; generative systems can also make a task newly accessible to someone who would otherwise have lacked the time, confidence, or specialist skill to attempt it. For an early-stage company, where the founder is frequently responsible for several functions at once, that expansion of capability matters as much as the reduction in time.

Noy and Zhang's experiment provides one of the clearest measurements of the effect. In a sample of **453** college-educated professionals completing occupation-specific writing tasks, access to ChatGPT reduced completion time by **40%** and increased average output quality by **18%** (Noy & Zhang, 2023). Writing may sound like a narrow task until it is placed inside the life of a young company. Proposals, websites, product descriptions, sales material, policies, support messages, investor documents, research summaries, job descriptions, and customer emails all require writing. Lowering the cost of each of those activities changes how much one person can produce before the company has generated enough revenue to hire specialists.

The same pattern appears in software development. Cui and colleagues ran three randomized field experiments across Microsoft, Accenture, and a Fortune 100 company, covering **4,867** developers. Access to an AI coding assistant increased completed tasks by **26.08%**, with particularly strong adoption and benefits among less experienced developers (Cui et al., 2026). For entrepreneurship, the significance is not that every founder becomes a software engineer. It is that the technical barrier to testing an idea, building a prototype, or working productively with developers becomes lower than it was when all meaningful implementation required scarce specialist labour.

Brynjolfsson, Li, and Raymond's study of **5,172** customer-support agents shows why the skill-distribution effect deserves attention. The AI assistant increased successfully resolved issues per hour by about **15%** on average, but newer and lower-skilled workers gained considerably more, approaching performance levels associated with much greater tenure (Brynjolfsson, Li, & Raymond, 2025). In other words, AI can compress part of the distance between an inexperienced worker and an experienced one. For a small firm, where hiring senior expertise in every function is rarely possible, that compression changes the minimum level of competence available to the organization.




Figure 1. Selected measured productivity gains from generative AI.

Sources: Noy & Zhang (2023); Cui et al. (2026); Brynjolfsson, Li & Raymond (2025). Outcome measures differ and should not be interpreted as directly comparable effect sizes.

The shrinking minimum viable company

The minimum viable company has been shrinking for years. U.S. Bureau of Labor Statistics data show that average employment at a startup fell from about 7.3 workers in 1998 to 3.5 in 2023, well before generative AI became widely available. Digital platforms, cloud software, outsourced services, online payments, and remote work had already reduced the amount of internal infrastructure required to start. Generative AI enters this trajectory by lowering the labour cost of several knowledge functions at the same time.

A very small team can now cover work that once required a designer, copywriter, analyst, junior developer, research assistant, or operations coordinator, at least at the level required to produce a first version. This should not be confused with eliminating the need for specialist expertise. Difficult engineering, sophisticated design, high-stakes analysis, legal work, and complex sales still require judgment and experience. The change is at the threshold: more tasks can reach an acceptable first version without a full specialist team, which means the company can move further before resource constraints force it to stop.

Business adoption data show that firms are already using AI in precisely these functions. The Federal Reserve's 2026 Small Business Credit Survey found current AI use among 46% of surveyed small employer firms, with another 15% planning adoption within a year. Among users, writing and marketing were by far the most common applications at 83%, followed by individual productivity at 61% and planning or analysis at 51%. Seventy-one percent of users reported productivity gains. These figures come from a weighted convenience sample rather than a random probability sample, so they should not be treated as a precise national prevalence estimate, but they are still useful for showing how strongly small-business adoption is concentrated in work that shapes the visible and administrative side of a company.




Figure 2. Most common AI uses among AI-using U.S. small employer firms.

Source: Federal Reserve Banks, 2026 Small Business Credit Survey. Weighted convenience sample.

Official data from other economies point in the same direction. Eurostat reported that 20% of EU enterprises with ten or more employees used AI technologies in 2025, up from 13.5% a year earlier. Statistics Canada reported business AI use of 19.2% in the second quarter of 2026, compared with 6.1% two years earlier. The UK Office for National Statistics reported that roughly a quarter of businesses were using some form of AI in late December 2025, while a separate government study produced a lower overall estimate but found the same concentration in language and text

generation. The percentages are not directly comparable because each survey defines its population and AI use differently, but the direction is consistent: AI capability is moving rapidly into ordinary business operations.




Figure 3. Selected official and institutional readings of current business AI use.

Sources: Federal Reserve Banks (2026), Eurostat (2025), Statistics Canada (2026), and UK Office for National Statistics (2026). Survey populations and definitions differ; percentages are directional rather than directly comparable.

The invisible effect of “good enough”

Measured productivity gains capture only part of the entrepreneurial effect because many early-stage decisions are not about doing a task faster; they are about whether the task is attempted at all. Before generative AI, a founder might decide against producing a market map because professional research was too expensive, postpone a prototype because development capacity was limited, or avoid testing ten versions of a sales message because creating them was too time-consuming. When the cost of a first attempt collapses, the decision changes from “Can we afford to do this?” to “Why not try it?”

That change expands the list of things a new business can produce before it has validated demand. It also raises the visual standard for what a launch is expected to contain. A polished website, coherent brand language, automated email sequence, clean proposal, and basic research pack once suggested that a company had invested substantial time and resources. Increasingly, these are accessible baseline outputs. The benefit is that smaller firms can present themselves professionally much earlier. The strategic consequence is that professional presentation carries less information about the depth of the organization behind it.

This is the first important shift in the report. AI makes competent production more accessible, but widespread access also changes what competent production signals. When many competitors can reach a similar minimum standard, the fact that a company looks polished becomes less useful as evidence that it understands the market better, has a stronger product, or is more likely to deliver. The company still needs those assets, but the advantage begins to move elsewhere.

CHAPTER 2

The Illusion of a Finished Business

A young company used to display its incompleteness openly. Early websites were rough, products were visibly partial, messaging changed constantly, and founders often knew exactly where the gaps were because those gaps required money or skills they did not yet possess. Generative AI can remove many of these visible signs before the underlying commercial questions have been answered. A new company can quickly acquire polished copy, automated communication, a coherent visual identity, a credible prototype, detailed market research, and a steady publishing rhythm. From the outside, it may look far more settled than the decisions underneath it actually are.

This matters because the unresolved questions in a new business are not primarily aesthetic. They concern whether the problem is important enough to change behaviour, whether the intended customer controls budget, what alternatives the buyer compares, why interested prospects fail to convert, which customers are expensive to serve, what causes retention, which channels can reach the audience repeatedly, and whether the company's own language matches the way customers describe the problem. None of those questions is answered simply because the surrounding material becomes more complete.

The **Production–Market Divide** is a way of separating these two forms of progress. Production progress occurs when the organization creates, improves, or automates something. Market progress occurs when uncertainty is reduced because the market responds. A new feature may represent both if customers use it and reveal something important; a revised pitch may do the same if it is tested in real sales conversations. The divide appears when the company continues to improve internal output without generating enough external evidence to know whether those improvements matter commercially.

Production progress and market progress

The distinction is easiest to see in ordinary operating decisions. A team can build five features around an untested assumption, publish a large volume of content without learning why prospects do not buy, or automate outbound activity before it knows which segment has urgency. It can optimize the conversion rate of a landing page while attracting an audience that is unlikely to purchase, or generate an elaborate customer persona without having heard the intended buyer explain the problem in their own words. None of this work is necessarily useless. The problem is that its value remains uncertain until it is linked to a market response.

A useful weekly discipline is therefore to separate two questions: **what did we create, and what did the market teach us?** The first answer may include a prototype, campaign, onboarding flow, revised pricing page, new CRM sequence, and a dozen pieces of content. The second answer should contain evidence that changed the company's understanding: perhaps procurement rejects the current contract structure, referrals close faster than paid leads, smaller customers are costly to serve, the “time-saving” message is weaker than a risk-reduction message, or a feature that users requested has little effect on retention. The value of the second list is not that it sounds more strategic; it is that it can make the company change course.

Output is not outcome

The risk is easy to state and difficult to manage because output is immediate while many outcomes arrive slowly. A landing page can be finished in a day, while a reliable answer about retention may require months. Positioning can be rewritten in minutes, while learning which message maps to a

real buying motive may require repeated sales conversations. Dashboards can be built quickly, whereas trust with a distribution partner may take a year. The production system therefore supplies fast, visible signals of progress, while the market supplies slower and often less flattering information.

This difference affects behaviour. Teams naturally gravitate toward work they can complete and measure. Generative AI strengthens that tendency because it keeps supplying additional tasks that feel productive: another variant, another analysis, another campaign, another automation. The issue is not that internal work should stop. Infrastructure, product quality, technical debt, and communication all matter. The issue is whether the organization has enough external learning to determine which internal work deserves priority.

The Progress Audit

EOS proposes a simple **Progress Audit** that separates activity into two ledgers. The first records production events: what the company created, changed, designed, coded, published, documented, or automated. The second records market-evidence events: information generated by customers, prospects, partners, channels, competitors, or observed behaviour that changed the firm's understanding of demand. The purpose is not to calculate a universal ratio. One enterprise contract or one decisive customer loss may carry more information than dozens of small production tasks. The purpose is to make the imbalance visible when one ledger is consistently full and the other remains vague.

A company can be very busy while learning very little. Generative AI makes that state easier to sustain because the supply of internal work is effectively unlimited. The discipline required in an AI-enabled firm is therefore not simply to produce faster, but to make sure that faster production remains tied to questions the market is actually in a position to answer.

CHAPTER 3

When Production Becomes Abundant, Something Else Becomes Scarce

A capability can remain useful while becoming less distinctive. The change occurs when access to that capability spreads faster than access to the other resources required for success. A business that once gained an advantage from professional design, competent copy, strong analysis, or basic software development may still need all four, but if competitors can obtain a similar level of quality through widely available AI tools, those capabilities become weaker sources of differentiation.

This is the **Entrepreneurial Scarcity Shift**. It does not claim that production no longer matters, or that every firm can produce equally well. It describes a change in relative value. When producing a credible first version becomes cheaper, the scarce resources move toward things that are harder to obtain on demand: sustained customer attention, trusted distribution, reputation, proprietary information, access to decision-makers, deep knowledge of a niche, recurring revenue, and the ability to learn from a market before competitors do.

From production scarcity to market scarcity

The distinction can be understood through a simple example. If a town has one excellent designer, professional design is scarce and can become a meaningful competitive advantage. If every local business gains access to tools that make competent design easy, the excellent designer may still be better, but the minimum standard of the market rises. Buyers become less impressed by the fact that a company looks professional because professionalism is no longer strong evidence of unusual capability. Their attention shifts toward other signals that remain difficult to reproduce.

A similar change may occur across several AI-assisted business functions at once. If most firms can generate competent copy, clean visuals, structured research, usable code, and polished presentations, then quality at the surface level rises across the field. The strategic question becomes less about whether the company can produce these outputs and more about whether it has access to information, relationships, customers, channels, and proprietary experience that competitors cannot generate from the same public tools.

Everyone gets better together

Doshi and Hauser's 2024 experiment provides a useful illustration of how this can happen. Participants who received AI-generated ideas produced stories that were rated as more novel; access to one idea increased novelty by about 5.4%, and access to five increased it by about 8.1%. At the same time, the total collection of AI-assisted stories became more similar. The study measured creative writing rather than commercial positioning or product strategy, so it should not be treated as direct evidence about startups. Its relevance lies in the pattern: a tool can improve individual performance while narrowing the distance between outputs produced by different people.




Figure 4. AI assistance increased individual novelty in Doshi and Hauser's creative-writing experiment.

Source: Doshi & Hauser (2024), Science Advances. The study also found that AI-assisted stories became more similar as a group.

Commercially, that possibility matters because many firms now ask similar models to solve similar problems using similar public information. They ask for homepage copy, positioning options, customer personas, content plans, outbound sequences, product names, or “premium” rewrites. The outputs may be considerably better than what an unassisted novice would have produced, but repeated exposure to the same conventions can also make brands, arguments, and product language feel familiar. The risk is not low quality; it is the gradual standardization of what competent work looks like.

The cost of looking professional falls

For many years, polished output acted as a rough signal of organizational depth. A strong website implied money and expertise; a professionally produced report suggested research capacity; a highly personalized proposal suggested that someone had spent time understanding the client. Those signals were never perfect, but they carried information because producing them required visible effort. AI lowers the cost of that effort, which is beneficial for access but changes how buyers interpret the result.

As presentation becomes cheaper, buyers may rely more heavily on evidence that remains costly to obtain: a respected referral, a track record, specific domain knowledge, customer advocacy, retention, the ability to answer detailed questions, or a reputation established over repeated interactions. These are not inherently “human” because they are sentimental or traditional. They are useful because they are harder to fabricate quickly and therefore carry more information about what the company has actually done.

Competence is not the same as differentiation

A common strategic mistake is to treat any quality improvement as a competitive advantage. If AI makes a company's copy **20%** better, that improvement is valuable. If it makes most competitors' copy better at roughly the same time, the improvement may be necessary simply to remain credible. Many technologies follow this path. Email, websites, search engines, cloud software, and spreadsheets became indispensable without remaining differentiators once they were widely available.

AI is likely to follow a similar pattern in at least some business functions. The more common the underlying models and interfaces become, the more durable differentiation will depend on what sits around them: proprietary customer data, relationships, operational knowledge, specialist judgment, original research, unique distribution arrangements, product history, and the quality of feedback flowing into the system. The key question is therefore not whether a company uses AI, but whether it combines broadly available production capability with market assets that are not broadly available.

CHAPTER 4

The Cliff Between Launch and Success

Lowering the cost of entry into entrepreneurship is valuable because it allows more people to test ideas, build independently, and enter markets that previously required more capital or a larger team. The difficulty is that easier entry does not automatically make survival or scale easier. These are different stages of the business process, governed by different constraints. AI is changing the first stage very quickly; the evidence reviewed here gives much less reason to believe that the later stages have been transformed to the same degree.

This distinction matters because the visible milestone of “launch” has moved earlier. A website can be live, the company can be registered, the initial product can function, and the brand can maintain a professional presence before the founder has accumulated much evidence about repeat demand. That does not make the launch false. It simply means that the point at which a company looks ready can now occur before many of the commercial questions that used to be exposed during a longer and more expensive build process.

The top of the funnel is enormous

U.S. Census Business Formation Statistics recorded **578,926** seasonally adjusted business applications in July 2026. The number conveys the scale of entrepreneurial intent, but its definition is important. A business application is an application for an Employer Identification Number; it is not proof that a company has begun trading, hired employees, earned revenue, or survived. For the July 2026 cohort, the Census projected that **29,959** employer businesses with payroll tax liabilities would form within four quarters. The two figures should not be converted into a simple “success rate,” because they measure different points in the statistical process, but they illustrate how much activity occurs before a business becomes an employing organization.

The filter becomes more severe when scale is considered over several years. OECD evidence on micro-startups finds that only about **4%** grow on average. In a stylized group of 100 micro-startups observed after five years, depending on the country, 26 to 58 are inactive, 36 to 71 remain below ten employees, and only one to eight reach ten or more employees. U.S. Bureau of Labor Statistics data tell a related survival story from a different angle: five-year survival rates across several historical startup cohorts have generally ranged from about one-half to the high-50% range, with the 2018 cohort at **57.3%**.




Figure 5. Five-year outcomes for a stylized cohort of 100 micro-startups across OECD countries.

Source: OECD DynEmp / Enabling SMEs to Scale Up. Ranges vary by country and predate the current generative-AI wave.

None of this should be interpreted as evidence that AI has made firms more likely to fail. Much of the business-demography evidence predates widespread generative AI and therefore describes a structural feature of entrepreneurship rather than a new AI-era outcome. Its relevance is comparative. If the cost and speed of producing a credible launch fall sharply while the historic difficulty of surviving and scaling remains, the relative weight of post-launch capabilities increases.




Figure 6. Five-year survival of selected U.S. startup cohorts.

Source: U.S. Bureau of Labor Statistics, Business Employment Dynamics.

Productivity gains and sales difficulty can coexist

The Federal Reserve's 2026 Small Business Credit Survey provides a useful snapshot of that coexistence. Among surveyed small employer firms, **46%** reported current AI use. Of the firms using AI, **71%** said productivity had increased, **39%** reported improved quality, and **31%** reported increased sales. These are meaningful benefits and they argue against any claim that AI is simply generating empty activity.




Figure 7. Reported productivity gains and customer-acquisition difficulty can coexist.

Source: Federal Reserve Banks, 2026 Small Business Credit Survey. The 71% figure is among AI-using firms; the 57% figure is among all employer-firm respondents. The juxtaposition is descriptive, not causal.

Yet the same survey found that 57% of firms viewed reaching customers and growing sales as an operating challenge, the most common challenge reported. Weak sales were a financial challenge for 48%. Among early-stage nonemployer firms planning to hire, 66% reported difficulty reaching customers or growing sales, while only 24% were profitable at the end of the referenced year and 93% expected revenue to grow over the following twelve months. The figures describe a business environment in which confidence, experimentation, and productivity can all be high while demand remains difficult to secure.

It is tempting to turn this juxtaposition into a stronger causal claim than the data allow. The Federal Reserve survey cannot tell us that AI caused customer-acquisition problems or that AI users are more likely to struggle with sales. Macroeconomic conditions, competition, rates, costs, industry structure, and business maturity all influence these outcomes. The defensible conclusion is narrower: rapid improvement in the tools used to produce business output does not eliminate the challenge of reaching a market, and the demand side remains a major source of difficulty even as AI adoption rises.

The wrong mental model of democratization

Public discussion of AI and entrepreneurship often moves too quickly from lower production costs to broader business success. The first part of the chain is convincing: if software, research, design, writing, and analysis become cheaper, more people can build and more people can launch. The mistake is assuming that the same change automatically creates more customer attention, more purchasing power, or more willingness to switch suppliers.

Consider an ecommerce business that can now launch at a fraction of the previous cost. The founder benefits from cheaper design, copy, photography, analysis, and automation. If thousands of other sellers receive the same benefit, however, the increase in supply may also intensify competition for attention. The tool has democratized the ability to enter the market without creating an equivalent increase in the amount of demand available to each entrant.

The same logic applies in software. A three-person team may be able to build what previously required a much larger engineering group, which is an extraordinary reduction in the technical barrier to entry. If many teams can now build similar products, the difficult question moves toward distribution, trust, integration into customer workflows, procurement, and retention. In professional services, AI can help small firms produce analysis, proposals, and outbound communication that look comparable to larger competitors, but this does not automatically provide the reputation or relationships that persuade a buyer to award the contract.

The point is not that competition necessarily becomes worse in every sector. It is that lower production barriers relocate the source of advantage. If more firms can reach the minimum standard required to compete, differentiation shifts toward the resources that remain unevenly distributed.

The launch is moving earlier in the learning process

Production friction historically forced some firms to learn from the market before they could afford to complete the product. Manufacturers sought orders before committing to large runs, consultants often needed a paying client before adding staff, and software founders frequently relied on design partners before building a broader team. This was not an ideal system; scarcity also prevented good ideas from being tested and excluded people who lacked capital. But the friction sometimes created a sequence in which customer evidence arrived early because production could not proceed indefinitely without it.

Generative AI makes it easier to reverse that sequence. A company can build, polish, automate, and publish before it has resolved basic questions about urgency, willingness to pay, or distribution. That sequence can still produce an excellent company, particularly when rapid prototyping is used to learn. The risk appears when cheap production becomes a reason to delay market exposure rather than accelerate it. In that case, the firm reaches launch earlier but reaches reliable market evidence later.

CHAPTER 5

The Optimization Trap

Internal business work has a psychological advantage over market work: it is more controllable. A founder can complete a homepage, revise a product, automate a process, create a campaign, or improve a presentation and end the day with visible evidence that something changed. A sales conversation may end with uncertainty, an interview may contradict the business plan, and a prospect may not respond at all. The market supplies information that the company cannot fully control, which is precisely what makes it valuable and what makes it easier to avoid.

Generative AI increases the supply of internal work. Once a company can create another version of a message, landing page, analysis, pricing model, persona, feature specification, or content plan in minutes, there is almost always something else to refine. The **Optimization Trap** occurs when this ability to improve internal output begins to compete with the need to expose the company to external judgment. The work may be useful and professionally executed, yet the organization can still become better at expressing or operating a weak assumption.

AI expands the safe side of entrepreneurship

Before generative AI, internal optimization encountered practical limits. Rewriting a website ten times, conducting repeated competitor analyses, mocking up many product concepts, or maintaining a large publishing schedule required enough labour that the team eventually had to prioritize. AI weakens those limits. The organization can continue producing variations at low cost, which is helpful when each variation is attached to a meaningful test but less helpful when variation becomes a substitute for testing.

This is why the problem should not be described as “founders are doing fake work.” The homepage may genuinely improve, the product may become easier to use, the CRM may become cleaner, and the content may be better. The issue is whether these improvements reduce commercially important uncertainty. A company that cannot explain why customers choose it, which segment has urgency, or what causes retention can become highly efficient at producing assets without becoming substantially more certain about its business model.

The Closed-Loop Founder

The **Closed-Loop Founder** is a useful way to describe what happens when the firm's information system begins with an assumption and repeatedly elaborates it without receiving enough external correction. The founder proposes a view of the customer, AI develops the view into personas and positioning, the company produces content or product changes based on that model, platform analytics measure the response, and AI then interprets the metrics and recommends another round of optimization. There may be real data in the process — traffic, clicks, open rates, usage, or engagement — but the system can still be weak if the underlying audience is poorly selected or the metrics sit too far from actual purchasing behaviour.

An open loop works differently because the market can change the premise rather than merely influence the execution. Customer behaviour creates observations; those observations alter the team's interpretation; decisions change; the company produces a new offer, feature, message, or process; and the next transaction attempt generates further evidence. AI can be present throughout this system, but its value comes from making the loop faster and more interpretable rather than replacing the external input.

The distinction becomes important when teams describe themselves as “data-driven.” A company can use large amounts of data and still optimize the wrong market, the wrong segment, or the

wrong behaviour. What matters is not simply the amount of information available, but whether the information is sufficiently close to the commercial decision the company is trying to improve.

The productivity paradox inside a startup

Productivity is usually measured within a defined process: more support issues resolved, more code completed, less time spent drafting, or more analysis produced. Those gains are valuable when the underlying task is already known to matter. A startup, however, is partly an execution system and partly a search process. It is still discovering which problem deserves attention, who should be served, what should be built, how the product should be priced, and which route to market can work repeatedly.

This creates a particular risk for AI-enabled teams. They can ship features faster before they know which problem carries urgency, produce more content before they understand who converts, or send more outbound messages before learning which accounts have a reason to change. The productivity gain is real at the task level, yet the firm may be accelerating activity that has not been commercially validated.

The better measure is therefore not simply whether AI increased output, but whether it increased the **rate at which the company removes important uncertainty**. A tool that allows a team to test a pricing hypothesis in two days instead of two weeks has a different strategic value from one that allows the same team to produce five times as much untested content. Both may count as productivity, but only one clearly tightens the market-learning loop.

Why the trap feels like strategy

Generative AI is particularly effective at producing coherent strategic language. A founder can request a go-to-market plan, market segmentation, competitive analysis, funnel, pricing framework, messaging architecture, risk register, or customer profile and receive a structured answer almost immediately. This is enormously useful for organizing thought, especially when the alternative is a blank page, but the speed and completeness of the output can create a sense of resolution that the underlying evidence does not justify.

A written strategy is not the same thing as a tested strategy. The quality of the articulation may improve substantially without changing the factual basis underneath it. The discipline required is to distinguish **strategic articulation** from **strategic evidence**. AI is capable of improving the first; the second still depends on data, transactions, customer behaviour, and observation.

The internal-output test

A simple diagnostic is to ask whether a recent period of work could have occurred almost unchanged if the company's prospective customers had disappeared. If a team can spend several weeks building, automating, publishing, and analyzing without receiving any information that could alter its commercial assumptions, the business may have a proximity problem rather than a productivity problem.

This does not mean that every week requires customer interviews or direct sales activity. Some products need long technical build cycles, infrastructure can be necessary, and established firms may have reliable market information embedded elsewhere in the organization. The warning sign is persistent separation: internal activity continues to increase while evidence about demand, purchasing, retention, or distribution remains weak. AI makes that condition easier to sustain because there is always more internal work available.

CHAPTER 6

The Market Cannot Be Simulated

AI is exceptionally good at producing plausible answers to questions such as who the ideal customer might be, what that customer cares about, what objections could arise, which price might work, or which channel appears suitable. These answers can accelerate preparation and expose blind spots, but they remain hypotheses until they are connected to observed behaviour. A synthetic persona does not control a budget, a generated objection does not delay a procurement process, and a predicted willingness to pay is not a transaction.

The distinction matters because simulated market knowledge can feel more complete than the evidence underneath it. A detailed persona may contain motivations, language, anxieties, purchasing criteria, and demographic context, all of which make it easier to reason about a market. Its usefulness depends on whether those assumptions are subsequently tested against people who can behave differently from the model. AI is most useful when it helps the company prepare better questions, organize observations, and compare evidence; it becomes risky when the plausibility of the generated answer is treated as evidence in itself.

The Market Evidence Ladder

EOS uses the **Market Evidence Ladder** to distinguish between different strengths of market information. The framework is not intended as a universal score, because buying cycles and forms of commitment vary widely across sectors. Its purpose is to remind teams that signals closer to costly behaviour generally deserve more weight than signals based on imagination or low-cost attention.

Level 0 — Imagined evidence

This includes founder intuition, AI-generated personas, hypothetical buyer profiles, competitor-inspired assumptions, and internal workshops. These inputs are useful for forming questions and designing tests, but they should remain clearly labelled as hypotheses rather than proof.

Level 1 — Attention evidence

Views, clicks, follows, opens, event attendance, and content engagement show that something attracted notice. Attention is necessary in many business models, but it is often cheap and can be generated by people who have little intention of buying.

Level 2 — Stated-interest evidence

Survey answers, interview enthusiasm, waitlist signups, demo requests, and statements such as “I would use this” move the firm closer to demand because the customer has expressed an intention. The limitation is familiar: stated preferences can change when price, effort, risk, procurement, or switching costs become real.

Level 3 — Friction evidence

Objections, negotiation, procurement questions, trial abandonment, delayed decisions, price resistance, and feature trade-offs are stronger signals because the customer has begun to confront an actual choice. Friction often reveals the part of the decision that early interviews miss.

Level 4 — Transaction evidence

A purchase, paid pilot, deposit, signed contract, or completed order is stronger because the customer has committed a scarce resource. It does not guarantee retention or long-term fit, but it establishes that demand moved beyond interest.

Level 5 — Retention and advocacy evidence

Repeat purchase, renewal, expansion, referral, recommendation, and unprompted advocacy show that value continued after the first transaction. In many business models, these are among the strongest available signals that the company has found something more durable than curiosity.

The ladder is useful because companies frequently compare evidence from different levels without naming the difference. “We have strong demand” can mean a large audience, a waitlist, a set of paid pilots, or hundreds of renewals. These are all useful observations, but they support very different levels of confidence.

Talking to customers is not enough

The corrective to over-simulation should not be reduced to generic advice to “talk to customers.” Song, Wang, and Parry studied **224** U.S. new ventures and found that formal collection and use of market information mattered for performance, while raw customer interaction did not have a significant direct relationship in their model (Song, Wang, & Parry, 2010). The distinction is important because interviews can become another form of activity that generates notes without changing decisions.

Customer contact becomes valuable when the company asks useful questions, samples the right people, interprets the evidence carefully, and allows the result to change priorities. AI can strengthen this process by structuring interview notes, comparing themes across conversations, identifying contradictions, and testing alternative interpretations. The system still depends on the quality of the external observations entering it.

Chen and Zhang’s 2024 research provides an additional boundary condition. Their multi-informant, three-wave study of **210** new ventures across Canada, Chile, and China found that moderate market-information acquisition was associated with better outcomes than either too little or excessive acquisition. Research itself consumes time. A firm can therefore replace one form of over-optimization with another if it keeps gathering information without deciding what to do with it.

The discipline of being wrong

Experimental research in entrepreneurship is valuable because it shows that learning is not merely about collecting positive confirmation. Camuffo and colleagues’ randomized controlled trial with **116** Italian startups taught one group to formulate and test business assumptions more scientifically. The treated founders performed better and changed direction more effectively. A later large-scale replication covering **759** firms across four randomized trials found that a scientific approach increased idea termination and search efficiency (Camuffo et al., 2020; 2024).

The importance of idea termination is easy to underestimate. Entrepreneurship culture often celebrates persistence, yet persistence in a weak model can consume years of capital and attention. Howell’s research on new ventures found that receiving negative feedback increased venture abandonment by around **13%** in the studied setting (Howell, 2021). The finding does not imply that negative feedback should always end a venture; it shows that useful information can create value by stopping activity as well as by encouraging it.

For an AI-enabled firm, this is particularly relevant because the cost of continuing to generate new versions has fallen. Cheap iteration is beneficial only when the organization retains the willingness to stop an idea that the market is repeatedly rejecting.

Faster building should mean faster testing

Koning, Hasan, and Chatterji's research on startup adoption of A/B testing found performance improvements on the order of 30–100% after one year, depending on the measure and specification. The value of experimentation was not simply that firms generated more options; it was that they could distinguish stronger ideas from weaker ones and scale the better-performing options more confidently.

Generative AI should make this logic more powerful because it reduces the cost of producing testable variants. If a firm can create ten landing pages, product messages, onboarding flows, or feature prototypes in the time it once took to create one, the limiting factor becomes the quality and speed of the test. Without a comparable improvement in experimentation, the organization simply accumulates more untested possibilities.

The wrong audience can make rigorous testing useless

Koning and Nanda's work on sampling bias shows why more experimentation is not automatically better. They studied a digital-product testing and launch environment in which roughly nine out of ten users were men. Products aimed at women experienced 45% less growth one year later when tested in that mismatched environment than products whose intended market more closely resembled the platform audience (Koning & Nanda, 2023). The problem was not a lack of data or a lack of experimentation. It was that the evidence came from the wrong people.

This is a useful warning for AI-era business systems because sophisticated tools can make weak samples look more authoritative. Dashboards, automated analysis, and statistically clean experiments do not protect a company from selecting an audience that does not represent the buyer it ultimately needs to serve. The relevant question is therefore not just how much data has been collected, but whose behaviour the data actually represents.

The strongest market signal is often inconvenient

External evidence earns its value partly by being capable of causing problems for the organization. It can reveal that the customer segment is wrong, the product is liked by users who cannot pay, the sales cycle is longer than expected, a feature has little effect on retention, or the obvious product benefit is not the actual buying motive. These findings are strategically useful because they constrain the company's freedom to tell itself a convenient story.

This is where market evidence differs most sharply from generated strategy. AI can help the organization interpret an uncomfortable result, identify alternative explanations, or design the next test, but it should not be used to explain away every contradiction. A learning system is valuable only if it retains the possibility that the current model is wrong.

CHAPTER 7

The Human Infrastructure of a Market

Markets are usually described through abstractions such as demand, supply, price, competition, and distribution. Those abstractions are necessary, but companies encounter them through people and institutions. Buyers exchange information with colleagues, suppliers develop reputations, procurement teams remember previous failures, customers recommend vendors, industry communities decide which providers are credible, and experienced operators know which introductions or channels are taken seriously. These relationships do not replace economic forces; they are among the mechanisms through which those forces operate.

Generative AI can map portions of this environment more efficiently than before. It can research accounts, identify likely decision-makers, summarize public discussions, prepare a salesperson for a meeting, and compare information across a large number of sources. What it cannot create instantly is the history that makes a market trust one participant more than another. That history is accumulated through repeated contact, delivery, referral, observation, and participation in the commercial system around the product.

What market embeddedness actually means

The term **market embeddedness** is useful only if it is defined in operational terms. A market-embedded firm has access to information that is difficult to infer from public sources alone. It knows how buyers describe the problem when they are not speaking in formal marketing language, which objections are genuine and which are polite ways to end a conversation, who signs a contract and who can quietly block it, what a credible vendor looks like inside the culture of the industry, and which events or communities actually influence purchasing behaviour.

This information is not necessarily secret. Much of it is simply accumulated through repeated exposure. A founder who has worked inside a sector for years may understand how budgets are approved, which regulatory changes create urgency, what buyers complain about privately, and why an apparently weaker competitor is trusted. AI can help organize and extend that knowledge, but the underlying advantage comes from the firm's position inside the market rather than from access to the model itself.

Networks are not a soft extra

Stam, Arzlanian, and Elfring's meta-analysis of 61 independent samples found a corrected correlation of .211 between entrepreneurs' network social capital and small-firm performance, with network diversity showing a particularly strong pattern in some contexts (Stam et al., 2014). The result should not be turned into simplistic advice that more networking automatically produces growth; the studies are observational and effects vary by firm age, industry, and institutional environment. What the evidence supports is that network resources can contribute economically meaningful information, introductions, legitimacy, advice, talent, capital, distribution, and customer access.

This helps explain why a founder with weaker technical resources can sometimes outperform a more technically sophisticated outsider. The outsider may understand the product more deeply, but the embedded founder may know which customer has a live budget, how the product fits into existing workflows, which partner can create access, or what evidence the market requires before taking a new vendor seriously. Production capability and market access are complementary, but they are not interchangeable.

The first customer is more than revenue

Denoo, Soh, and Clarysse's 2026 research on first-customer acquisition among technology ventures identified multiple routes to the first customer, including strong existing ties, weak ties, and market mechanisms that did not depend on prior relationships. Their findings caution against the idea that existing connections are the only route to early sales. A venture can win a first customer without a pre-existing relationship, and in their sample different acquisition paths were associated with different later outcomes.

The larger point is that the first customer changes the information available to the company. Before a real customer exists, many beliefs about value remain hypothetical. Once a customer begins using the product or service, the firm encounters onboarding friction, unexpected use cases, price sensitivity, support demands, procurement issues, and the customer's own language for the problem. Revenue matters, but the first transaction also gives the organization a live environment in which its assumptions can be tested.

Relationship quality affects performance

Palmatier and colleagues' meta-analysis of relationship marketing found that relationship quality and relationship investment can materially influence performance, with effects especially relevant in B2B, services, and channel settings where personal relationships often carry substantial weight (Palmatier et al., 2006). In these contexts, trust and credibility can be transferred through people who already have standing inside the market.

As generated communication becomes easier, this transfer of trust may become more important as a signal. Buyers can now receive more polished outreach from more firms, and many providers can present themselves competently online. A trusted introduction, known track record, or respected customer reference is harder to reproduce on demand because it depends on prior behaviour. The point is not that human relationships are inherently superior to technology; it is that some of the information a buyer needs is carried by reputation rather than by the quality of the message alone.

Selling is an adaptive act

Sales research offers one of the clearest examples of why market capability cannot be reduced to asset production. Franke and Park's meta-analysis covered **155** samples and more than **31,000** salespeople and found that adaptive selling was positively related to multiple measures of performance. Goad and Jaramillo's later synthesis, covering approximately **126,790** salesperson responses, similarly highlighted the importance of adapting behaviour to the customer.

A sales script can now be generated almost instantly, but selling still involves interpretation. The seller must distinguish a real price objection from a trust problem, recognize whether the decision-maker is present, decide whether the buyer is asking for more information or trying to exit politely, and understand whether the customer values speed, risk reduction, compliance, status, convenience, or another motive that may not have appeared in the original pitch. AI can support these judgments by preparing the seller and surfacing patterns across calls, but the commercial advantage comes from responsiveness rather than from the existence of a well-written script.

The information inside relationships

One of the most valuable functions of market relationships is that they provide information before it becomes formal. People inside an industry often know that a regulation is becoming painful, a supplier is losing trust, a major customer is replacing a system, a competitor has disappointed key accounts, or a new budget is becoming available before those changes appear in public data. A company that participates deeply in the market can hear these signals early and use AI to analyze them alongside broader information.

This suggests a more durable model of AI advantage. As foundation models become broadly available, the model itself may be difficult to defend as a proprietary asset. The quality of the information entering the model can be much harder to copy. Private customer conversations, sales outcomes, support history, product usage, operational knowledge, partner relationships, and original research give the company inputs that a competitor cannot obtain from the same public prompt.

CHAPTER 8

The Authenticity Problem

The falling cost of communication changes the signals carried by communication. A personalized email once implied that someone had spent time preparing it; a long-form article suggested research effort; polished creative work implied access to production capacity; a detailed customer response suggested that a person had paid attention to the case. Those assumptions were never completely reliable, but they were reasonable because producing the work required time and labour. Generative AI weakens the connection between the apparent effort in a message and the actual effort required to create it.

This does not mean that customers will reject AI-generated content or that human authorship is always preferred. It means that the receiver has less reason to treat the existence of polished communication as proof of expertise, care, or organizational depth. When a signal becomes cheaper to produce, buyers look for other evidence that remains difficult to manufacture quickly.

When effort becomes difficult to see

The effect is especially important in markets where trust, judgment, or personal attention are part of what is being sold. A consultancy can claim expertise, a service provider can promise responsiveness, and a founder can say that the company understands the customer's industry, but the buyer must still decide how seriously to take those claims. Historically, a deeply researched proposal or highly specific piece of advice carried some credibility because producing it required knowledge and effort. AI reduces that cost and therefore changes what the output proves by itself.

More durable signals include a visible track record, customer advocacy, retention, specific domain knowledge, respected referrals, performance history, and the ability to answer detailed questions when the interaction cannot be pre-scripted. These signals can still be exaggerated or manipulated, but they are harder to generate instantly because they depend on prior experience and other people's behaviour.

What the experiments tell us

Research on AI-authored marketing is beginning to identify where the trust boundary lies. Kirk and Givi conducted seven preregistered experiments examining responses to messages believed to have been written by AI or humans. When emotional marketing messages were perceived as AI-authored, positive word-of-mouth and loyalty fell, with perceived inauthenticity and moral disgust helping explain the effect. Importantly, the penalty weakened in several conditions: when messages were factual rather than emotional, when AI edited human work rather than fully replacing it, when AI openly represented itself or the brand, and when consumers believed AI use was already common.

Brüns and Meißner found a related pattern across three experiments involving social content. GenAI-created content could reduce perceived brand authenticity and produce less favourable responses, but the penalty was smaller when AI assisted humans rather than replacing them. Grigsby, Michelsen, and Zamudio similarly found that AI disclosure in service advertising could reduce trust and advertising attitudes, especially when the message emphasized intangible or human qualities of the provider. Dai and colleagues provide an important counterpoint: consumers can prefer AI for objective information while preferring humans for more subjective information, and satisfaction improves when the type of agent fits the type of information being provided.

Taken together, these studies do not support the broad claim that customers distrust AI. They suggest that the customer's expectation of the interaction matters. Where the communication is valued for speed, consistency, retrieval, or routine factual explanation, AI can be entirely appropriate and may be preferred. Where the communication is supposed to demonstrate empathy, judgment, responsibility, personal knowledge, or emotional sincerity, replacing the human signal can carry a cost.

The human layer should be designed, not preserved by habit

The practical conclusion is not to keep people involved in every customer interaction. That would sacrifice legitimate productivity gains and preserve expensive work that customers may not value. A more useful approach is to map interactions according to **trust sensitivity**. Routine status updates, document summaries, invoice information, and standard comparisons are different from negotiations, high-risk recommendations, apologies after a serious failure, or decisions in which a customer needs to know who is accountable.

The design question is therefore whether human presence changes the customer's interpretation of the interaction. A buyer may appreciate an AI-generated technical comparison and still want a human expert to recommend which option carries the least risk. A customer may welcome automated shipping updates but react differently to an automated apology after a major service failure. The same technology can be appropriate in one part of the customer journey and damaging in another.

The content abundance problem

A related issue is the growth in the supply of content. When every firm can create more material at lower cost, the amount competing for a fixed amount of human attention can increase quickly. The available evidence does not yet justify a macro-level claim that generative AI has caused an economy-wide collapse in engagement or a universal rise in customer-acquisition cost. Those remain research questions rather than established facts.

The logic of abundance is nevertheless important for strategy. If attention is scarce, responding to weaker engagement by producing larger quantities of generic material may treat an attention problem as a production problem. Content is more likely to remain valuable when it contains something competitors cannot reproduce from the same public inputs: original evidence, proprietary data, a customer case, specialist experience, a useful tool, a credible community relationship, or an argument grounded in years of operating inside a particular market.

AI can help research, structure, and present those assets. What it cannot supply automatically is the history that gives them weight. That history becomes increasingly important as the surface quality of generated communication rises.

CHAPTER 9

AI Is Not the Problem

A report about the risks of AI-enabled overproduction can easily slide into an argument against the technology itself. That would misread the evidence. AI can improve customer service, raise productivity, support analysis, help less experienced workers perform at a higher level, and contribute to sales as well as efficiency. Qualitative research on startup growth hacking also shows firms using generative AI for market research, market-entry analysis, customer engagement, personalization, and promotional work. The important distinction is therefore not whether a company uses AI, but where AI sits inside the company's information and decision process.

The same tool can either insulate the organization from the market or help it become more responsive to the market. That difference depends on what the AI is being asked to replace. When it replaces the need to gather evidence, it can create confidence without validation. When it reduces the cost of analyzing evidence, testing alternatives, or preparing for higher-quality customer interactions, it strengthens the market-learning loop.

AI used to avoid the market

AI becomes strategically weak when generated material is allowed to stand in for information the firm should be collecting directly. A synthetic persona may replace customer interviews, predicted objections may replace listening to real sales calls, generated customer language may substitute for the words buyers actually use, and a model-generated market map may become the basis for a go-to-market plan that has never been tested through access to the segment it describes. In each case, the company receives a coherent representation of the market without having to experience the friction that would reveal whether the representation is correct.

The danger is not that the AI output is always wrong. It may be informed, useful, and directionally accurate. The problem is that its confidence can exceed the evidentiary basis available to the firm. When internally generated analysis is followed by internally generated execution, the organization can spend substantial time refining a model that has never faced strong external correction.

AI used to shorten the distance to the market

The same technology can be used in the opposite way. Customer interviews can be compared for recurring themes, sales-call notes can be searched for repeated objections, support tickets can be clustered into emerging problems, and churned customers can be compared with retained ones before the team decides whom to call. Account research can improve the quality of a sales conversation, test variants can be produced faster, and administrative work can be removed from senior people so that more of their time is available for customers and partners.

These uses do not replace market evidence; they increase the organization's capacity to absorb and act on it. This is the more useful interpretation of AI as a **market amplifier**. The technology creates leverage around the relationship, experiment, or decision rather than pretending that the external input is no longer necessary.

The strongest AI systems may be fed by proprietary reality

As foundation models become widely available, access to the model itself becomes less likely to provide a durable advantage. What one firm can ask a public system, another firm can often ask as

well. Differentiation therefore shifts toward the information and operating context that surround the model.

Private customer conversations, product-usage histories, support patterns, sales outcomes, operational knowledge, partner relationships, original research, and long-term customer data are much harder for competitors to reproduce. A firm that has spent years in a market can use AI to search, compare, and interpret that proprietary evidence at a scale that was previously difficult, while a new entrant using the same base model lacks the underlying information.

This changes the strategic question. Instead of asking only how much AI can be added to the business, management should ask how much **better reality** can be added to the AI system. The model is useful because it can work across a larger body of evidence; the defensible advantage lies in owning or accessing evidence that others do not have.

AI should increase exposure, not insulation

A practical test of an AI deployment is therefore behavioural. After adoption, does the team reach customers more effectively, analyze feedback faster, run better experiments, identify weak ideas sooner, prepare more intelligently for sales conversations, and surface useful information from support or churn? Or does the organization mainly produce more documents, messages, campaigns, and analyses without increasing its exposure to the market?

This standard is more demanding than counting AI use cases because it connects the tool to commercial learning rather than to activity. An AI system that saves ten hours of administrative work has created capacity; whether that capacity becomes strategically valuable depends on what the company does with the time it recovers.

CHAPTER 10

The Market-Embedded AI Company

The practical challenge is not to choose between technological leverage and human connection. A well-designed AI-native firm should be able to use automation aggressively while preserving the external feedback loops that keep it attached to demand. EOS describes this operating model as the **Market-Embedded AI Company**: an organization that uses AI to reduce the cost of production, analysis, repetition, and synthesis while deliberately measuring whether those gains are being converted into better market information and stronger commercial evidence.

The framework is organized around six dimensions. It is not a predictive score and should not be treated as one. Its value lies in helping management locate the point at which a fast-moving company has become strong at producing but weak at learning.

Dimension 1 — Production velocity

Question: How quickly can the company turn an idea into something that can be tested?

AI can improve production velocity dramatically, but the relevant goal is not maximum output. The goal is to reduce the time between a meaningful hypothesis and a usable test. A team should be able to create a prototype, message, offer, workflow, or analysis quickly enough that it does not spend weeks debating an assumption that could be exposed to evidence much sooner. Production velocity becomes counterproductive when the team starts measuring success by the number of assets created rather than by the uncertainty those assets help remove.

Management prompt: What did we build this month that became a real test, and what did we build that simply became more internal inventory?

Dimension 2 — Customer-contact velocity

Question: How frequently does the organization encounter behaviour from real customers or prospects?

Customer contact includes more than interviews. Sales conversations, onboarding, service delivery, support tickets, procurement, renewals, cancellations, referrals, product usage, partner discussions, and field observation can all produce market evidence. What matters is that the signal originates outside the company's own assumptions. A firm with high production velocity and low customer-contact velocity is vulnerable because its ability to create changes faster than its ability to learn whether those changes matter.

Management prompt: When was the last time a customer or prospect told us something important enough to change what we planned to do?

Dimension 3 — Experiment velocity

Question: How quickly are the company's most important assumptions forced to meet evidence?

Not every decision requires an experiment, but assumptions that could materially damage the business if false deserve early testing. These include whether the customer actually has the problem, whether the problem is urgent, whether the intended buyer controls budget, whether the price is acceptable, whether a channel can acquire customers economically, and whether a product change influences willingness to pay or retention. AI can lower the cost of creating tests, but the commercial benefit appears only when the test reaches the right audience and the result is allowed to affect the next decision.

Management prompt: Which belief are we currently spending the most money or time on without having seriously tried to disprove it?

Dimension 4 — Learning conversion

Question: Does market information actually change decisions?

Companies can accumulate large quantities of feedback without becoming more market-driven. Interview transcripts, survey results, dashboards, support data, and AI summaries may all increase while the roadmap, pricing, segment, channel mix, and sales strategy remain unchanged. **Learning conversion** asks whether new evidence enters the operating system of the company. A useful measure is not how much research was collected, but which decisions would have been different if the research had never existed.

Management prompt: Name three material decisions in the last quarter that changed because of customer or market evidence.

Dimension 5 — Relationship density and distribution access

Question: How deeply connected is the company to the people and channels through which the market actually moves?

Audience size is not the same as relationship density. A firm can have a large online following and weak access to buyers, while another may know a relatively small number of people in a niche and possess unusually strong routes to information, referrals, partnerships, and distribution. The relevant pattern depends on the business model: consumer products may rely more heavily on product-led or platform distribution, enterprise services may depend on trusted introductions, local businesses may rely on reputation and repeat customers, and regulated sectors may require institutional credibility.

Management prompt: If paid advertising disappeared tomorrow, which customers, communities, partners, or channels could still carry the company into the market?

Dimension 6 — Revenue evidence

Question: What has the customer actually committed?

The Market Evidence Ladder is most useful when the organization is explicit about the strongest behaviour supporting an important growth assumption. Attention, stated interest, demos, transactions, renewals, and referrals all carry information, but they are not equivalent. A company with 100,000 views and a company with 500 renewals may both describe “strong demand,” yet the evidence behind the statement is fundamentally different.

Management prompt: What is the strongest customer behaviour supporting our most important belief about future growth?

Taken together, the six dimensions create a different way of evaluating an AI-enabled company. Production remains important, but it is judged by how efficiently it creates testable work. Customer contact is judged by the quality of external signals. Experimentation is judged by whether the right assumptions are being challenged. Research is judged by whether it changes decisions, relationships by whether they create real access, and demand by the strength of customer commitment. The result is a management system that treats market learning as an operating capability rather than as a periodic research exercise.

CHAPTER 11

The EOS Operating System for AI-Native Entrepreneurship

The six dimensions become more useful when they are translated into operating habits. The aim of the EOS system is not to slow an AI-enabled company down; it is to prevent production velocity from moving so far ahead of market learning that the organization becomes efficient at serving an assumption. The practices below are designed for founders and small teams that can now create far more material and functionality than traditional startup processes assumed.

1. Start every build cycle with a market uncertainty

Teams should define the uncertainty before defining the asset. “We need a new landing page” is a production request; “we do not know whether mid-sized firms respond more strongly to compliance risk or labour savings” is a market question. Once the uncertainty is clear, the landing page becomes an instrument for testing rather than an end in itself. The same principle applies to content, product features, pricing pages, demos, and outbound campaigns.

This habit changes the review conversation. The team no longer asks only whether the asset was completed or whether it looks good. It asks whether the asset produced evidence capable of changing a decision.

2. Separate generated insight from observed evidence

Strategy documents should distinguish claims based on internal assumptions, AI synthesis, secondary research, customer statements, observed behaviour, and transactions. The categories should not be treated as equivalent simply because they appear together in a polished document. A useful internal discipline is to label statements according to whether the company **thinks**, **observed**, **tested**, or **knows** something. The purpose is not bureaucratic precision; it is to prevent generated plausibility from acquiring the status of market evidence through repetition.

3. Give AI-produced strategy an external test

When AI recommends a segment, the next question is whether the company can reach that segment and whether the segment responds. When it proposes a message, the message should be used in a real buying context. When it suggests pricing, a transaction should eventually test the price. When it predicts an objection, the sales team should listen for whether customers actually raise it. The aim is not to prove the model wrong but to keep strategic recommendations connected to behaviour that exists outside the model.

4. Build an external information moat

The company should deliberately accumulate information that competitors cannot obtain from the same public system. Structured sales notes, interview archives, churn explanations, proprietary surveys, support histories, account maps, product behaviour, partner relationships, field observations, and original research all create a body of evidence that becomes more valuable as general AI capability spreads. The organization should also make this evidence reusable by storing it in forms that can be compared over time rather than leaving it fragmented across inboxes, individual memories, and disconnected software.

5. Automate around the relationship rather than through it

Automation is most valuable when it removes friction surrounding high-value human interactions. AI can prepare a call, summarize an account, draft a follow-up, log notes, compare conversations, surface risks, and schedule the next action. If the relationship itself carries trust, accountability, judgment, or market information, replacing it entirely may destroy part of the value the system is supposed to support. The decision to automate should therefore consider not only task cost but what the customer or partner is meant to infer from the interaction.

6. Review market evidence with the same seriousness as product output

Most product companies have regular reviews of what was built. Fewer organizations review what the market taught them with the same discipline. A monthly market-learning review can examine the assumptions held at the beginning of the period, the evidence that arrived, which beliefs became stronger or weaker, which activities stopped, and which questions now require testing. This creates institutional memory and reduces the chance that a rejected assumption returns later in a new piece of generated strategy.

7. Reward work that is correctly stopped

If experimentation is functioning, some projects should end. Organizations usually reward completed work and visible output, which creates a bias toward continuation. AI makes completion cheaper and therefore makes that bias more dangerous: a weak idea can be given better copy, more features, better visuals, and more analysis without becoming a better business. Teams should recognize the value of discovering early that an idea, segment, feature, or channel does not justify further investment.

8. Track the Production–Market Divide

The Progress Audit can be used as a recurring management check. The team reviews what it produced and separately reviews what the market taught it. No universal ratio is required, because businesses vary in their build cycles and sales models. The purpose is to identify persistent divergence. If production volume rises rapidly after AI adoption while market evidence, customer contact, and transaction learning remain flat, management should investigate whether the new productivity is being converted into commercial learning.

A Reader's Diagnostic — Are You Building Faster Than You Are Learning?

The following diagnostic is a discussion tool rather than a validated score. It is intended to help a management team examine the last sixty to ninety days of activity and identify where the information loop may be closing in on itself.

1. We can point to many new assets, features, campaigns, or automations, but fewer concrete changes in what we know about customers.
2. Our team uses AI-generated personas or market analysis more often than direct customer language.
3. We can produce new marketing material faster than we can explain why customers actually choose us.

4. Engagement metrics receive more attention than purchasing, retention, referral, or sales friction.
5. Important assumptions are being funded or built around without a transaction-level test.
6. Customer feedback is collected but rarely kills projects or changes priorities.
7. AI has increased output more clearly than it has increased customer contact or experimentation.
8. Customer-facing interactions are automated without distinguishing routine tasks from trust-sensitive ones.
9. If paid distribution disappeared, our direct routes into the target market would be weak.
10. Competitors can access most of the same information we use to make strategic decisions.
11. We spend substantial time polishing work that could already be tested.
12. We know more about what customers say they want than what they repeatedly pay for.
13. Our strongest proof of demand remains close to attention or stated interest rather than transaction, retention, or referral.
14. We have not deliberately built a proprietary archive of customer, sales, product, or market learning.
15. The company could continue its current weekly routine for a long period without having a difficult conversation with a real buyer.

The useful outcome is not the number of statements that apply. The diagnostic is most valuable when it helps the team identify which part of the commercial learning system is weak: exposure to customers, quality of testing, interpretation of evidence, conversion of learning into decisions, distribution access, or strength of revenue proof.

CHAPTER 12

The New Entrepreneurial Advantage

The founder introduced at the beginning of this report can still create an extraordinary amount of work between Friday evening and Sunday night. That capability should not be minimized. A small team can reach a level of technical and presentational competence that once required more people, more capital, or more time, and this creates opportunities for founders who previously would have been blocked by gaps in skill or resources.

What changes is the strategic meaning of that capability as it becomes common. A finished-looking website, functioning prototype, professional proposal, research summary, and content operation are easier to produce than they were, which means their presence alone tells the market less about the company behind them. The more useful question is no longer how much one person can build, but how quickly the organization can use what it builds to obtain better information about demand and convert that information into a stronger position.

From leverage to learning

The first phase of AI adoption has naturally focused on leverage. Organizations can do more with fewer people, reduce time spent on repetitive work, and give less experienced workers access to capabilities that were previously harder to obtain. Those gains are likely to remain important and in many sectors will become part of the expected operating baseline.

The competitive difference will depend on what companies do with the capacity they recover. If every saved hour is reinvested in more internal production, the organization may simply create a larger volume of work. If some of that capacity is directed toward sales conversations, better experiments, customer research, account development, support analysis, distribution partnerships, and original evidence, the same productivity gain strengthens the parts of the business that are harder for competitors to copy.

The moat after the model

Foundation models, interfaces, and agentic tools will continue to improve and become more accessible. As that happens, capabilities that are available to nearly every firm become less reliable as sources of sustained advantage. This does not make them unimportant; it means they become part of the infrastructure required to compete.

Durable advantage is more likely to sit in assets that remain unevenly distributed: customer history, reputation, relationships, proprietary data, judgment about a niche, trusted distribution, knowledge of why particular deals were won or lost, original research, and the organizational willingness to change course when evidence is unfavourable. AI can make these assets more useful by increasing the amount of information the company can search, compare, and interpret. It does not remove the need to accumulate them.

The human company after AI

The most AI-enabled firms may therefore become more deliberate about where they preserve human involvement. This is not an argument for artificial informality, founder-centric branding, or a vague preference for “authenticity.” It is a commercial design question. Human presence matters most where trust, responsibility, judgment, negotiation, and high-quality market information are

created through interaction. AI can remove work around those moments without automatically replacing the moments themselves.

The same principle applies internally. A company should not judge an AI system only by how much output it generates or how much labour it saves. It should also ask whether the system improves the quality of decisions by connecting the organization more closely to evidence that originates outside itself.

The new scarcity

The evidence reviewed in this report supports a clear change on the production side of entrepreneurship. Writing can be faster and better, coding assistance can increase completed work, customer-support productivity can rise, and AI adoption is spreading quickly across firms and countries. It also shows that reaching customers and growing sales remain major operating challenges, that scaling from a micro-startup remains rare, and that experimentation, market learning, networks, relationship quality, and adaptive selling continue to matter for performance.

Those findings do not establish that AI has increased startup failure or made customer acquisition more expensive. They support a more precise conclusion: AI has reduced the cost of assembling the supply side of a business faster than it has reduced the difficulty of earning a durable position on the demand side. That difference changes where management attention should go.

The opportunity for founders is to use AI not simply to produce more, but to shorten the distance between an assumption and the evidence that can confirm or reject it. The companies most likely to build a defensible advantage will combine widely available production tools with market assets that remain difficult to generate on command: customers who return, relationships that create access, data accumulated through real use, distribution that can be repeated, and the knowledge that comes from being deeply involved in the market the company is trying to serve.

Research Notes and Evidence Boundaries

What the evidence supports strongly

The evidence reviewed for this report strongly supports the following propositions:

- Generative AI can reduce time and raise output quality in several forms of knowledge work.
- AI can increase productivity among software developers and customer-support workers.
- Less experienced workers can receive disproportionately large productivity benefits in some settings.
- Business AI adoption is rising across several advanced economies.
- Writing, marketing, productivity, analysis, and language-based work are major business use cases.
- Customer acquisition and sales growth remain major challenges for small businesses.
- Scaling from micro-startup to larger employer is historically rare.
- Structured experimentation and scientific decision-making can improve entrepreneurial search.
- Entrepreneurial social capital, relationship quality, market orientation, and adaptive selling are meaningfully related to firm or salesperson performance.
- AI assistance can increase individual creative output while reducing diversity across a collection of outputs.
- AI authorship can create authenticity or trust penalties in some customer-facing contexts.

What the evidence does not yet establish

The evidence reviewed does not establish that:

- AI has caused startup failure rates to increase.
- AI has caused customer-acquisition costs to increase across the economy.
- AI-created startups fail more often than non-AI-created startups.
- AI use causes founders to spend less time with customers.
- AI has caused a measurable economy-wide collapse in differentiation.
- AI content is generally ineffective.
- Customers generally distrust AI.
- Greater AI use is negatively associated with startup success.

These remain possible research questions rather than established conclusions.

Important survey caveats

The 2026 Federal Reserve Small Business Credit Survey is a weighted national convenience sample rather than a random probability sample.

Official AI-adoption surveys differ significantly in wording, population, firm-size thresholds, and definition of AI. Adoption percentages should therefore be compared as directional evidence, not treated as measurements of one identical construct.

Business applications in U.S. Census Business Formation Statistics are applications for Employer Identification Numbers, not operating firms or successful startups.

OECD micro-startup scale-up figures describe structural business-demography patterns and should not be interpreted as evidence of an AI-era causal change.

Core References

- Barari, M., et al. (2021). A meta-analysis of customer engagement behaviour. *International Journal of Consumer Studies*. <https://doi.org/10.1111/ijcs.12609>
- Brüns, J. D., & Meißner, M. (2024). Do you create your content yourself? Using generative artificial intelligence for social media content creation diminishes perceived brand authenticity. *Journal of Retailing and Consumer Services*. <https://doi.org/10.1016/j.jretconser.2024.103790>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at Work. *The Quarterly Journal of Economics*. <https://doi.org/10.1093/qje/qjae044>
- Camuffo, A., et al. (2020). A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial. *Management Science*. <https://doi.org/10.1287/mnsc.2018.3249>
- Camuffo, A., et al. (2024). A scientific approach to entrepreneurial decision-making: Large-scale replication and extension. *Strategic Management Journal*. <https://doi.org/10.1002/smj.3580>
- Chen, J., & Zhang, Y. (2024). Too much of two good things: The curvilinear effects of self-efficacy and market validation in new ventures. *Journal of Business Research*. <https://doi.org/10.1016/j.jbusres.2024.114845>
- Cui, Z., et al. (2026). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Management Science*. <https://doi.org/10.1287/mnsc.2025.00535>
- Dai, et al. (2026). AI or human: How the type of information to be disclosed alters customer service agent preferences. *Journal of Retailing and Consumer Services*. <https://doi.org/10.1016/j.jretconser.2025.104621>
- Denoo, L., Soh, P.-H., & Clarysse, B. (2026). From strong ties to no ties: Configurations for first-customer acquisition in tech startups. *Journal of Business Venturing*, 41(2), 106570. <https://doi.org/10.1016/j.jbusvent.2025.106570>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*. <https://doi.org/10.1126/sciadv.adn5290>
- Ellis, P. D. (2006). Market orientation and performance: A meta-analysis and cross-national comparisons. *Journal of Management Studies*. <https://doi.org/10.1111/j.1467-6486.2006.00630.x>
- Franke, G. R., & Park, J.-E. (2006). Salesperson Adaptive Selling Behavior and Customer Orientation: A Meta-Analysis. *Journal of Marketing Research*. <https://doi.org/10.1509/jmkr.43.4.693>
- Goad, E. A., & Jaramillo, F. (2014). The good, the bad and the effective: A meta-analytic examination of selling orientation and customer orientation on sales performance. *Journal of Personal Selling & Sales Management*. <https://doi.org/10.1080/08853134.2014.899471>
- Grigsby, J. L., Michelsen, M., & Zamudio, C. (2025). Service ads in the era of generative AI: Disclosures, trust, and intangibility. *Journal of Retailing and Consumer Services*. <https://doi.org/10.1016/j.jretconser.2025.104231>
Corrigendum: <https://doi.org/10.1016/j.jretconser.2025.104427>
- Howell, S. T. (2021). Learning from Feedback: Evidence from New Ventures. *Review of Finance*. <https://doi.org/10.1093/rof/rfab006>
- Kirk, C. P., & Givi, J. (2025). The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications. *Journal of Business Research*. <https://doi.org/10.1016/j.jbusres.2024.114984>
- Koning, R., Hasan, S., & Chatterji, A. (2022). Experimentation and Start-up Performance: Evidence from A/B Testing. *Management Science*. <https://doi.org/10.1287/mnsc.2021.4209>
- Koning, R., & Nanda, R. (2023). Sampling Bias in Entrepreneurial Experiments. *Management Science*. <https://doi.org/10.1287/mnsc.2021.01740>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Palmatier, R. W., Dant, R. P., Grewal, D., & Evans, K. R. (2006). Factors Influencing the Effectiveness of Relationship Marketing: A Meta-Analysis. *Journal of Marketing*, 70(4), 136–153. <https://doi.org/10.1509/jmkg.70.4.136>
- Read, S., Song, M., & Smit, W. (2009). A meta-analytic review of effectuation and venture performance. *Journal of Business Venturing*. <https://doi.org/10.1016/j.jbusvent.2008.02.005>

- Rezazadeh, et al. (2025). Generative AI for growth hacking. *Journal of Business Research*. <https://doi.org/10.1016/j.jbusres.2025.115320>
- Sinha, et al. (2023). Meta-analytic evidence on entrepreneurial orientation, learning orientation and firm performance. *Journal of Business Venturing Insights*. <https://doi.org/10.1016/j.jbvi.2023.e00415>
- Song, M., Wang, T., & Parry, M. E. (2010). Do market information processes improve new venture performance? *Journal of Business Venturing*, 25(6), 556–568. <https://doi.org/10.1016/j.jbusvent.2009.03.003>
- Stam, W., Arzlanian, S., & Elfring, T. (2014). Social capital of entrepreneurs and small firm performance: A meta-analysis of contextual and methodological moderators. *Journal of Business Venturing*, 29(1), 152–173. <https://doi.org/10.1016/j.jbusvent.2013.01.002>
- Vaccaro, M., Almaatouq, A., & Malone, T. W. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-024-02024-1>

Official and Institutional Data Sources

- Federal Reserve Banks.** (2026). *2026 Report on Employer Firms: Findings from the 2025 Small Business Credit Survey*. <https://www.fedsmbusiness.org/2026-report-on-employer-firms>
<https://doi.org/10.55350/sbcs-20260303>
- Federal Reserve Banks.** (2025). *2025 Report on Nonemployer Firms: Findings on Hiring Plans from the 2024 Small Business Credit Survey*. <https://www.fedsmbusiness.org/reports/survey/2025/2025-report-on-nonemployer-firms>
<https://doi.org/10.55350/sbcs-20251014>
- U.S. Census Bureau.** (2026). *Business Formation Statistics — July 2026*. <https://www.census.gov/econ/bfs/current/index.html>
- U.S. Census Bureau.** (2026). Business Trends and Outlook Survey reporting on AI use. <https://www.census.gov/library/stories/2026/05/ai-use-businesses.html>
- U.S. Bureau of Labor Statistics.** (2024). *Business Employment Dynamics Twentieth Anniversary!* <https://www.bls.gov/spotlight/2024/business-employment-dynamics-twentieth-anniversary/>
- OECD.** *Enabling SMEs to Scale Up: Strengthening SMEs and Entrepreneurship for Productivity and Inclusive Growth*. https://www.oecd.org/en/publications/strengthening-smes-and-entrepreneurship-for-productivity-and-inclusive-growth_c19b6f97-en/full-report/component-5.html
- Eurostat.** (2025). *20% of EU enterprises use AI technologies*. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2>
- Statistics Canada.** (2026). *Analysis on artificial intelligence use by businesses in Canada, second quarter of 2026*. <https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2026010-eng.htm>
- Office for National Statistics, United Kingdom.** (2026). *Business insights and impact on the UK economy — 8 January 2026*. <https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/bulletins/businessinsightsandimpactontheukeconomy/8january2026>
- UK Department for Science, Innovation and Technology.** (2026). *AI Adoption Research*. <https://www.gov.uk/government/publications/ai-adoption-research/ai-adoption-research>
- World Bank.** *Entrepreneurship Database*. Formal limited-liability company registration data should be treated as formal-sector entrepreneurship measures and not as a complete measure of informal entrepreneurial activity.