

Timeless startup principles for founders

Nine principles for better decisions, from the first idea to scale and exit

Research-informed founder guide · Revised 23 September 2026

A startup makes consequential decisions before it has reliable answers. Your job is to find an opportunity worth pursuing, turn uncertainty into useful evidence and build a business that can deliver repeatedly. Eventually, that means committing decisively while uncertainty still remains.

Build something customers choose. Find what limits progress now. Match the next commitment to the evidence and the cost of waiting.

These nine principles address different ways a startup can get stuck. Use the one that changes the decision in front of you; keep the others as checks on what that decision could damage.

Principle	The question it answers
1 Customer choice	Why will customers change what they do?
2 Decisive learning	What must we discover before the next commitment?
3 Focus	Where can we build a position worth having?
4 Adoption	How does interest become a realised customer outcome?
5 Economics	What does a sale earn, cost and consume in cash?
6 Scaling	What deserves more resources now?
7 Ownership	Who can deliver and decide without a founder rescue?
8 Advantage	What becomes stronger with each cycle?
9 Freedom to act	Can we still act if our preferred route is delayed?

Start with a real choice: a hire, pilot, product investment, customer contract, production line, new market or funding round. Write down the commitment, identify the constraint, and decide what would justify proceeding, changing the plan or stopping.

The next page shows how. Each principle then provides an operating rule, evidence to seek and a boundary against misuse. The guide also includes evidence by business type, stage transitions, difficult trade-offs, two worked decisions and an editable decision record.

The principles are a synthesis of research and entrepreneurial practice. They improve the discipline of decisions; they do not guarantee success or eliminate the need for founder judgement. The evidence and its limits appear at the end.

Questions or feedback during the pilot?

Guido van Gageldonk | Venture developer, BOM

[Contact Guido via the BOM AI-tech page](#)

Use the principles to make one decision

You do not need nine parallel workstreams. Start with one consequential decision and the constraint that most limits progress. A constraint is the thing whose improvement would unlock the next meaningful outcome: demand, technical feasibility, adoption, delivery capacity, economics or cash.

1 Name the commitment and the alternative

Write: "We are deciding whether to ... by ...". Include the money, founder attention, people, capacity and obligations involved. State the best alternative use of those resources. Include the cost of waiting: lost customers, delayed learning or a closing market window.

2 Separate what you know from what you need to believe

List the few assumptions that could change this decision. Label each as observed, inferred or unknown. Name the strongest evidence against your preferred option. One exceptional customer, a successful demo and a funding offer support different claims.

3 Find the smallest investigation that can change the choice

Choose evidence relevant to the next commitment. A customer conversation may resolve a buying question; it cannot establish production yield. A pilot line may be the smallest credible technical test. Set an owner, cost limit, time limit and a decision date.

Before starting, complete three sentences: **proceed if ...; change the plan if ...; stop this route if ...**. Derive the criteria from the customer outcome, business economics and consequences of being wrong. The example thresholds in this guide are illustrations, not startup benchmarks.

4 Decide and commit

When evidence is sufficient for the proposed step, take it. Further analysis has a cost. When evidence contradicts a critical assumption, change the route or stop that commitment. When the result is inconclusive, name the remaining uncertainty and explicitly authorise a bounded follow-up; do not quietly call the test a success.

Keep the decision, rationale, owner and review trigger in one place. Write down what you expect to observe next. This makes judgement inspectable without pretending that you can predict everything.

A practical founder rhythm

For short learning cycles, a weekly discussion can cover four things: the last customer outcome, the current constraint, the new evidence and the next commitment. Review cash and longer-term dependencies on a cadence that matches their speed of change. Deeptech experiments may take months; commitments and emerging signals still need attention between milestones.

Use the decision record on page 17. If the discussion ends with a longer list of activities but no changed decision, priority or explicit commitment, the principles have not yet done their job.

1 Make customers choose you

Deliver an outcome valuable enough to change real behaviour, and verify that customers actually receive it.

A customer compares your offer with the current alternative, including doing nothing. The relevant improvement includes switching effort, learning, integration, risk and ongoing use. A superior feature can still lose that comparison. In consumer markets, the outcome may be enjoyment or identity; it need not be relief from a painful problem.

Use this principle when: choosing a problem, assessing a pilot, prioritising a feature or deciding whether interest is strong enough to justify more investment.

What to do

- Specify the customer, situation and outcome. Describe how the job is done today and why a change matters now. In B2B, identify the user, buyer, budget owner and person who can block adoption.
- Observe a recent real example before asking for opinions. Then ask for a commitment appropriate to the stage: access to a workflow, scarce engineering time, a purchase, a deployment or a repeat order.
- Define success before delivery and check afterwards. Measure the complete outcome, including work shifted to the customer. Record why customers decline, abandon or fail to return.

Evidence that changes the decision

Proceed: relevant customers make meaningful commitments, achieve the intended outcome and show reasons to keep choosing the offer. Repeated independent examples strengthen the case; the required evidence depends on the size of your next commitment.

Change the plan: users benefit but the buyer has no budget, the improvement disappears during implementation, or the value only exists for one unusual customer. Revisit the segment, scope, price or buying proposition.

Stop this route: well-targeted attempts repeatedly fail to find an outcome or commitment that could support the intended business, and you have no credible new explanation to test.

Where founders misapply it

A paid pilot may buy exploration rather than operational adoption. Agree who could approve normal use and what must happen first. For a one-off machine sale, follow-on use and field performance may matter more than monthly retention. Customer requests reveal needs; founders still design the solution.

Blank's search framework and Graham's early-customer practice support this approach. They are practitioner guidance, rather than causal estimates of success. [\[11\]](#)[\[12\]](#)

Next decision: identify the next customer commitment that would materially strengthen your case, and ask for it.

2 Attack the assumption that matters most

Spend learning effort on the uncertainty whose resolution could change your next important decision.

Founders need conviction about a direction and honesty about the route. Without a testable claim, almost any result can justify continuing. Equally, endless testing can postpone a decision for which you already have enough evidence. Useful learning changes what you will do.

Use this principle when: the team disagrees about a bet, a roadmap expands, a result is being presented as validation or further research keeps delaying action.

What to do

- List what must be true for the next commitment to work. Include customer behaviour, technical performance, access, cost and timing. Prioritise by decision impact, uncertainty and the cost of obtaining useful evidence.
- State the claim and an alternative explanation. "Users like our demo" could mean they value the outcome, enjoy novelty or want access to your team. Design the test to distinguish explanations that matter.
- Write the expected result, acceptance criteria and action branches before observing the outcome. Use realistic conditions, retain failures and ask someone to challenge the interpretation.

Evidence that changes the decision

Proceed: the critical claim survives a relevant test strongly enough to justify the proposed commitment. State what remains unproven; passing one test does not validate the whole company.

Change the plan: the outcome is mixed or the test did not measure the claim. Repair the test, narrow the use case or investigate an alternative explanation. Give the follow-up its own budget and deadline.

Stop this route: a necessary assumption fails a credible test and no affordable, plausible change can repair it. Sunk effort is not new evidence.

Where founders misapply it

Do not optimise easy metrics while avoiding a difficult existential question. In small B2B samples, precise percentages can disguise weak evidence. Combine observed behaviour, technical results and explicit commitments without pretending they are a large statistical sample. Some questions require expensive trials; "smallest useful test" does not mean "cheapest activity".

Camuffo et al.'s randomised studies support more disciplined search and idea termination. Research on A/B testing informs digital experimentation; transfer to other settings requires judgement. Neither validates every template in this guide. [\[1\]\[3\]](#)

Next decision: complete: "If we learned that ..., we would change ...". If the second blank stays empty, reconsider the test.

3 Focus where you can win

Concentrate resources on a reachable market where you can create distinctive value and a business worth building.

A large problem is not automatically an accessible opportunity. You need a customer group you can reach, a reason it will choose you and an economic opportunity consistent with your ambition. Market selection can matter more than working harder on the current plan.

Focus makes learning transferable: similar customers help product, sales and delivery improve together. A narrow entry point can support a large company if it opens a credible expansion path. A profitable niche can also be the intended destination; the financing model must fit that choice.

Use this principle when: choosing an initial segment, accepting a large unusual customer or considering another sector, country or product line.

What to do

- Compare a few plausible entry markets on customer urgency, access, feasibility, competition, economics and timing. A fatal weakness cannot be averaged away by a high total score.
- Choose the primary customer and use case. State which attractive requests you will decline and show where team capacity actually goes. A strategy with no refusals has not allocated scarce resources.
- For your intended scale, explain the route beyond the entry market. Identify what transfers: distribution, product, qualification, reputation or delivery. List what must be proved again.

Evidence that changes the decision

Proceed: you can reach suitable customers, their responses support the proposition and serving them improves a common core. The opportunity fits the founders' goals and the capital being raised.

Change the plan: every sale requires a different product or buying process, the timing depends on developments beyond your reach, or the entry market cannot support the intended outcome. Explore an alternative with bounded capacity.

Stop this route: there is no credible way to reach, serve or earn enough from the market within the resources you can assemble.

Where founders misapply it

Focus is a current allocation choice, not a promise to ignore evidence. Revisit it when the underlying assumptions change. Domain experience can improve access and judgement, but familiarity is not proof of opportunity. Research associates relevant industry experience with success; it does not require a particular founder age or background. [\[4\]](#)

Porter's trade-offs and activity fit explain the strategic mechanism. Ecosystem research also shows why the readiness of complementary technologies matters. [\[10\]](#)[\[8\]](#)

Next decision: name one opportunity you will refuse and one observation that would justify revisiting that refusal.

4 Build the path from interest to adoption

Make it possible for the right customer to find you, buy, start successfully and realise the promised value.

Distribution and implementation belong in the business design. They are not tasks to discover after the product works. An offer can fail at procurement, trust, integration, training or a partner hand-off even when users want it.

Use this principle when: demonstrations go well but sales stall, pilots do not convert, onboarding needs founder rescues or a partnership is expected to produce growth.

What to do

- Trace one real customer from first contact to a useful outcome and continued use. For each step, record the decision-maker, action, delay and reason to stop. Locate the step where progress actually breaks down.
- Learn the initial selling process yourself or alongside the person doing it. Test whether acquisition works beyond warm introductions. Separate a customer's willingness to buy from your ability to reach similar customers economically.
- Include onboarding, qualification, integration and support in the offer. Give each essential step an owner. For partners, establish the incentive, capacity and specific activity that should produce customer progress.

Evidence that changes the decision

Proceed: customers move through the path, reach the intended result and do so with increasingly understood effort. You can explain which acquisition route produces suitable customers and why.

Change the plan: interest is real but one step repeatedly fails. Fix that step before adding more leads. The answer may be a smaller initial deployment, a different buyer, a clearer contract or implementation capacity.

Stop this route: an essential channel, approval or complement remains unavailable and no workable alternative fits the business. A partner logo does not overcome that dependency.

Where founders misapply it

Founders should not wait for a perfectly scalable channel before serving early customers. Personal acquisition and manual onboarding can reveal what the eventual system needs. Record the work so it is not mistaken for effortless product adoption. Own the customer outcome; you do not have to own every company or asset involved in delivering it.

Graham explains the value of early manual work. Teece explains how control of complementary assets, including distribution, affects who captures innovation's returns. [\[12\]](#)[\[9\]](#)

Next decision: choose the single broken step in the buying or adoption path, assign an owner and test a repair with a real customer.

5 Know what each sale earns and consumes

Understand the revenue, delivery effort, acquisition cost and cash timing behind growth.

Creating customer value, retaining an economic share and receiving cash in time are separate achievements. Growth can strengthen the business or amplify a loss. A profitable contract can still create a cash crisis when suppliers and employees must be paid before the customer pays.

Use this principle when: setting price, accepting custom work, subsidising adoption, hiring to deliver demand or raising capital against a growth plan.

What to do

- Choose the right unit: customer, contract, transaction, installed machine or licence. Record revenue and attributable delivery costs, including implementation, service, infrastructure, quality failures and recurring human effort.
- Show acquisition separately and make founder effort visible. Distinguish one-off learning costs from work that recurs for every customer. Explain the specific mechanism expected to improve the economics.
- Map when cash leaves and arrives under a base case and a downside case. Include working capital and capacity investments. Test price, scope and payment structure as seriously as product features.

Evidence that changes the decision

Proceed: the next commitment has a credible economic purpose: an attractive contribution, a justified learning investment or a demonstrated path towards better economics. The cash required can be funded.

Change the plan: sales grow but delivery effort, cash exposure or customer concentration grows faster. Reprice, narrow the offer, redesign delivery or change payment terms before multiplying the pattern.

Stop this route: the viable price and achievable cost do not support the intended model, and no credible improvement or strategic reason justifies continuing the subsidy.

Where founders misapply it

This principle does not demand immediate profitability. Deeptech development, network creation and product discovery can require substantial investment before attractive unit economics are observable. Make the investment thesis, milestones and funding dependency explicit. "Scale will fix it" needs a mechanism and evidence.

A positive customer contribution still has to support R&D, overheads and investment. Lifetime-value estimates are fragile when history is short. Track observed contribution and cash alongside longer-term assumptions. Skok's definitions are useful; SaaS ratios are not universal thresholds for every business. [\[14\]](#)

Next decision: calculate what one representative sale earns and consumes today, then identify the assumption that makes the growth plan economically credible.

6 Move fast; scale on evidence

Act quickly on reversible choices and increase fixed commitments when the evidence justifies the exposure.

Early scaling magnifies untested assumptions. Excessive hesitation also destroys opportunities. The question is not whether to be bold or cautious: it is which part of the business has earned a larger commitment, what uncertainty remains and what waiting would cost.

Use this principle when: hiring ahead of revenue, expanding sales, increasing production, entering a new market or debating whether to accelerate after traction.

What to do

- Separate what works from what is still being discovered. Evidence for one segment, channel or production batch does not automatically transfer to another.
- Identify the current constraint. If qualified demand exceeds reliable delivery capacity, improve delivery. If interest does not become adoption, more salespeople may amplify the wrong problem.
- Choose the size and reversibility of the next bet. Define what extra capacity should produce, what must not deteriorate and when you will review the result. Include the cost of delay in the decision.

Evidence that changes the decision

Proceed and accelerate: relevant customer outcomes repeat, the delivery and economic logic support expansion, and the additional resources relieve a demonstrated constraint. Complete certainty is not required.

Change the plan: the new customer cohort, location or production run behaves differently. Keep the proven core moving while investigating what no longer transfers. Limit the experiment's claim on the core business.

Pause or stop the expansion: extra capacity fails to improve the intended outcome, quality or economics deteriorate beyond agreed limits, or the funding needed to complete the bet is no longer credible.

Where founders misapply it

Some evidence requires scale. A pilot line, regulatory study or locally dense marketplace may be the minimum useful commitment. Define why that scale is necessary and which assumptions it will test. Conversely, do not hide indecision behind another discovery sprint when customers and delivery already justify investment.

Lee and Kim find an association between early scaling and higher failure, using job postings as a scaling proxy. This supports scrutiny of timing; it does not establish a universal hiring limit or prove that early hiring causes failure. [\[2\]](#)

Next decision: state exactly what is being scaled, what earned that decision and what observation would make you change its size or pace.

7 Build a team that owns outcomes

Put capable people in charge of clear outcomes, with the information and authority to act.

A startup initially relies on founders to connect product, customer and resources. As it grows, that concentration can become a queue. Hiring does not solve the problem if every important decision still waits for the same person or if bad news arrives too late.

Use this principle when: choosing a co-founder or early hire, repeatedly rescuing delivery, struggling with conflicting decisions or finding that people avoid reporting problems.

What to do

- Start with the outcome the business needs and the capability missing today. Assess candidates using relevant work and evidence, not prestige alone. Discuss ambition, roles and how difficult disagreements will be resolved.
- Give an important outcome one accountable owner. Agree decision rights, resources, quality boundaries and when to escalate. Transfer context and customer understanding as well as tasks.
- Make problems discussable. Ask what failed, what was learned and what changes next. Combine candour with clear expectations; tolerating repeated avoidable underperformance is not psychological safety.

Evidence that changes the decision

Proceed: a colleague or team repeatedly makes sound decisions and delivers the outcome within agreed boundaries. Widen responsibility and remove approvals that no longer protect anything useful.

Change the plan: decisions stall or results require constant rescue. Diagnose whether the gap is skill, context, authority, incentives or capacity before assuming another hire will fix it.

Stop the arrangement: an unresolved role conflict, persistent lack of capability or concealment of material problems prevents reliable execution. Address the role or relationship directly rather than letting the whole organisation compensate indefinitely.

Where founders misapply it

Delegation does not remove founder responsibility for direction, key people and material commitments. Hands-on involvement remains useful when learning is essential. Repeated rescue should teach the system something. Chronic overload is also a capacity and continuity problem; build backup for critical knowledge and decisions.

Research on early employees highlights their lasting organisational importance. Psychological safety research supports learning behaviour, while management experiments show the potential value of basic operating discipline in their studied settings. These findings do not validate an “A-player” hiring formula. [\[5\]\[6\]\[7\]](#)

Next decision: transfer one recurring outcome with its decision rights and boundaries, then check the result without taking every decision back.

8 Build advantages that compound

Use each cycle of selling, building and delivering to create assets that improve future customer value or economics.

Revenue is valuable, but two companies with the same revenue can be building very different futures. One accumulates reusable capability; the other accumulates obligations. Ask what remains after the work is finished and why that makes the next cycle better.

Use this principle when: choosing between customer projects, prioritising platform work, evaluating an intellectual-property strategy or claiming that growth creates a defensible position.

What to do

- Name the asset and the mechanism. Production learning might improve yield; an evaluation system might reduce failure; a distribution relationship might lower acquisition effort. A patent, dataset or customer count is not its own explanation.
- Track a consequence that matters: customer outcome, reliability, time to value, cost, willingness to pay or access. Check whether the effect improves across comparable cycles.
- Establish what you control and can reuse. Identify rights, dependencies and who gains bargaining power as the business grows. Reassess whether a technical or market change could erode the advantage.

Evidence that changes the decision

Proceed: the work creates a reusable asset and the expected benefit appears in later customers, batches or releases. Invest in the mechanism that makes the improvement repeat.

Change the plan: every contract adds maintenance without strengthening a common core, more data does not improve outcomes, or a partner captures most of the benefit. Change scope, architecture, rights or the route to market.

Stop the claim: you cannot explain or observe how the supposed advantage improves value or economics. Keep it as an unproven hypothesis rather than using it to justify more spending.

Where founders misapply it

Do not spend years constructing a moat around demand that has not been established. Early custom work can be valuable when it earns money, teaches something important or opens access; record its full cost and what can be reused. Network effects need evidence that additional participation improves value for other participants. User growth alone does not demonstrate that mechanism.

Teece distinguishes creating innovation from capturing its returns and explains the role of complementary assets and imitation. The particular advantage still has to be demonstrated in your company. [\[9\]](#)

Next decision: complete: "After the next set of customers or production cycles, we will be better at ... because ..., and we will observe it in ...".

9 Preserve the ability to choose

Keep enough time, trust and control over critical dependencies to act when the preferred plan changes.

Keeping options open forever is not a strategy. Building a company requires commitment. The aim is to make that commitment without allowing one delayed payment, funding round, supplier or founder to decide the company's future by default.

Use this principle when: raising capital, taking on a dominant customer, entering an exclusive partnership, making a long-term commitment or considering an acquisition or IPO.

What to do

- Start from the next meaningful evidence or commercial milestone. Estimate the time and cash to reach it, the time to secure the next funding source and a reasoned allowance for setbacks. Identify when a fallback must begin to remain workable.
- Discuss the desired company and outcome among founders and capital providers. Compare financing choices on growth expectations, future funding needs, control and who receives value in different outcomes.
- Keep essential rights, contracts, financial information and responsibilities understandable. Document decisions as they happen. Be explicit about product limitations and material deviations from promises.

Evidence that changes the decision

Proceed: the commitment has a credible route to completion and a downside the company can bear. Critical counterparties can deliver, and the founders understand the dependencies they are accepting.

Change the plan: the milestone slips, funding takes longer, concentration increases or a partner becomes unreliable. Act while a smaller plan, different funding route, partnership or sale remains possible.

Stop the commitment: it would consume the remaining ability to respond without a credible path to a valuable outcome. Renegotiate, reduce the scope or change the route before choice disappears.

Where founders misapply it

Cash conservation alone cannot rescue a venture that never tests its central claim. Nor is dependence always avoidable: young deeptech companies may rationally rely on external capital or specialised suppliers. Make the dependence explicit and monitor the event that would make it untenable.

Sarasvathy's affordable-loss approach and Graham's funding-dependence question inform this principle. An acquisition or IPO adds its own readiness requirements; neither is the compulsory destination of every startup. [\[15\]\[13\]\[16\]\[17\]](#)

Next decision: identify the earliest trigger that would require a different plan, and the latest date at which that alternative can still work.

Use evidence that fits your business

The principles stay stable; the observable proof changes. Select evidence for the claim and commitment at hand. There is no universal revenue threshold, number of customers or product-market-fit score that validates every startup.

B2B software and AI

Look for an agreed workflow outcome, a real buyer, successful implementation and continued use or renewal. Track acquisition, onboarding, support and any recurring human effort. For AI, test accepted work under representative conditions, including checking and correction; a persuasive demonstration does not establish reliable operation.

Typical false positive: a paid innovation pilot with heavy founder support, no operational budget and no decision-maker for deployment. **Next proof:** a route to normal buying and a comparable customer outcome with understood delivery effort.

Industrial hardware and deeptech

Separate technical feasibility, qualification, manufacturability, field performance and commercial demand. Examine yield, tolerances, supplier lead times, service obligations and cash tied up before acceptance. Customer engineering time, a qualification programme or a conditional order can be meaningful early evidence, while leaving later risks unresolved.

Typical false positive: a working prototype interpreted as proof of economical production. **Next proof:** the particular test batch, pilot line or field trial that resolves the uncertainty behind the next capital commitment. Ecosystem readiness may determine whether a good technology can be adopted. [\[8\]](#)

Consumer products and marketplaces

Study repeat behaviour among comparable users and whether acquisition brings people who continue to get value. For a marketplace, examine whether both sides can complete valuable transactions in the relevant locality or category. Separate genuine repeated value from subsidies, novelty and founder intervention.

Typical false positive: sign-ups, downloads or transaction volume without evidence of sustained value and plausible economics. **Next proof:** a bounded market in which people return and successful transactions occur under conditions you could afford to reproduce.

Science-based and regulated ventures

Technical milestones, validation pathways, adoption and financing may move on different clocks. Map which development or approval step unlocks which next choice. Commercial exploration can continue while technical work proceeds, without treating customer interest as proof of efficacy, safety or approval.

Typical false positive: progress on one development axis presented as progress on every axis. **Next proof:** the evidence required for the specific technical, commercial or financing decision. The applicable validation and regulatory standards remain specific to the product and market.

Across all four: write down what the evidence establishes, what it leaves open and what commitment it is strong enough to support.

Let evidence define the next stage

Funding rounds and company age are useful context. They do not establish readiness. A company can scale a proven core while searching in a new market. The founder's operating task changes as the dominant uncertainty changes.

From infancy to a credible opportunity

Clarify the customer context, the technical possibility, the founders' ambition and the resources available. Identify what could make the direction untenable. The next commitment should buy meaningful evidence, such as access to a real workflow or a feasibility test, rather than merely a more polished pitch.

Transition question: is there a coherent opportunity worth investigating further, and a workable way to investigate it?

From first outcomes to product-market fit

Find a recognisable customer group that chooses the offer and receives meaningful value. Investigate adoption, willingness to pay and repeated behaviour using measures appropriate to the product. Distinguish a valuable solution from one customer's willingness to sponsor development.

Transition question: do the outcomes and commitments form a repeatable pattern, and what important condition would break that pattern?

From founder-led delivery to repeatability

Have another person or team reproduce a material part of selling, implementation or delivery. Understand effort, quality and economics; retain founder involvement where it is still needed for learning. Choose capacity additions around the actual constraint.

Transition question: what now works with understandable effort, and which larger commitment does that justify?

From repeatability to scale

Expand what is working while tracking whether new cohorts, locations or batches retain its strengths. Give teams clear outcomes and decision rights. Build processes when they prevent recurring errors or delays. Investigate what each expansion changes instead of assuming that all previous proof transfers.

Transition question: is growth strengthening customer outcomes, delivery and economic value, or hiding deterioration in averages?

From operating value to an ownership transition

A buyer needs to understand what transfers: customers, team, rights, technology and delivery capability. Consider the value it can create with the business and the founders' alternatives. An IPO requires continuing public-company capabilities in reporting, governance and leadership, as well as an attractive financing rationale. [\[16\]](#)[\[17\]](#)

Transition question: does this route improve the likely outcome versus realistic alternatives, and can the company keep operating if the transaction is delayed? A sale may make sense before independent scale; an IPO is not a final certificate of success.

Resolve the trade-off in front of you

Principles are most useful when two good arguments conflict. Name the relevant condition, then choose. These decision rules are a practical synthesis; they are not universal empirical thresholds.

Tension	Decision rule
Persist or stop	Continue while new evidence supports a viable route. Change or stop when a necessary assumption fails and there is no credible repair worth funding. Do not confuse difficulty executing with evidence that the opportunity is absent.
Move fast or test more	Move now when the choice is reversible and delay costs more than additional information is worth. Investigate more when the commitment is hard to reverse or the consequences of error are severe.
Focus or explore	Keep the primary commitment clear. Give an adjacent opportunity a bounded test when its potential justifies the distraction. Expansion earns resources through evidence, not novelty.
Customise or standardise	Accept custom work when its contribution, learning or access justifies all delivery and future maintenance costs. State what strengthens the reusable core and what is a deliberate exception.
Hire or fix the system	Hire when a capability or capacity gap is the constraint. First change the process or decision rights when capable people are blocked by recurring ambiguity.
Conserve cash or invest	Fund the evidence or capacity that can change the company's position. Size the commitment against available resources and delay costs. Spending little on inconclusive work can still be expensive.
Follow customers or lead	Use customers to understand context, outcomes and adoption. Retain responsibility for the solution and strategic choices. A requested feature is one possible route to an outcome.

A disagreement protocol for founders

Ask each person to state the decision they support, the assumption behind it and what evidence would change their mind. Identify whether the disagreement concerns facts, risk tolerance or the company you want to build. More customer interviews will not resolve different ambitions between co-founders.

Use the agreed decision owner and boundaries. Where you can test the disputed assumption at reasonable cost, do so. Where you cannot, make a judgement, record it and set a review trigger. Reopening the same decision without new evidence consumes capacity too.

A useful final check: are we choosing this because it improves the business, or because it protects a past decision, an investor narrative or our identity as founders?

Worked decision: an industrial AI startup

This fictional case shows how the principles change a real choice. All amounts, sample sizes and time limits are illustrative, not recommended benchmarks.

A team helps engineers assess and document machine-design changes. Two firms like the demo; a third offers a paid pilot. The founders propose six hires and expansion into three industries. The decision is whether to create that fixed cost now.

The critical claim and the next test

The team narrows the first use case. It must reduce the total time to an accepted change assessment, including checking and rework. Engineering users, a quality representative and the budget owner agree the baseline and acceptance criteria.

They authorise an eight-week, €40,000 pilot using representative tasks, with human approval retained. For this customer, a 25% reduction in total time at the agreed quality would justify considering deployment. Any uncontained critical failure pauses the trial. Passing this limited sample would not establish reliability for all tasks or uses.

The team also records who can authorise regular buying, the implementation work required and what a second comparable deployment must demonstrate. This separates technical usefulness from adoption and repeatability.

The economics behind the proposal

Assume a €30,000 annual contract, €4,000 in model and hosting costs and €18,000 in direct implementation, support and review effort. The provisional contribution is **€8,000 before acquisition, R&D and overheads**. The team must distinguish first-customer learning from costs that will recur. A low model bill alone does not justify scaling.

Agree the action branches before testing

Proceed to a larger commercial commitment: the customer outcome meets the agreed criteria, an operational buyer commits and the delivery effort supports a credible offer. Begin a comparable second deployment and add the capacity that this work actually needs.

Change the plan: quality or time savings depend on an overly broad scope, or value exists without a buying route. Narrow the task or resolve the budget and implementation gap with a bounded follow-up.

Stop this use case: no affordable revision can meet the required outcome, or the attainable price cannot support a credible delivery model.

The observed result and the founder decision

In this example, the time and quality criteria are met, but the pilot sponsor cannot secure operational budget. The team postpones the six-hire expansion. It spends the next commercial effort on a budget-owner decision and a second comparable customer; it reviews which implementation cost is reusable.

The pilot produced useful evidence. It did not prove that a broad sales organisation was the next investment. The founder decision changes because the remaining constraint is now visible.

Worked decision: it is time to accelerate

This fictional software case tests the opposite failure mode: treating uncertainty as a reason to wait after the business has earned a larger commitment. The observations describe this example, not a universal scale-up threshold.

A workflow product serves one defined B2B segment. Successive customer groups achieve the agreed outcome and continue using and paying for the product over the relevant observation period. A salesperson has won suitable customers beyond the founders' network. The team understands acquisition and delivery effort, and current contracts make a positive contribution before central costs.

Demand now waits because onboarding requires specialist capacity. Customers face delayed start dates. The founders consider another broad customer-discovery exercise, hiring more salespeople or increasing implementation capacity.

Identify what is already supported

The evidence supports demand and delivery for this segment and offer. It does not establish lifetime retention, success in a second industry or performance at ten times the volume. Those uncertainties matter, but they do not all need resolving before removing today's onboarding constraint.

Waiting has a visible cost: delayed customer value, postponed cash receipts and a risk of losing committed customers. More leads would worsen the queue. General discovery would add little to the current decision.

Make the bounded scaling commitment

The founders approve an experienced implementation hire and automate the most repeated setup task. A named owner is responsible for reducing the onboarding queue while maintaining agreed customer outcomes. The cash forecast includes recruitment, salary and training before any expected capacity benefit.

Before approval, they define the target improvement from the queue and workload data, a quality floor, a spending limit and a review after enough implementations to assess the change. They also schedule an earlier check for failed onboarding or unexpected cash pressure. The commitment is specific; it is not permission for a general hiring wave.

Decide what the next evidence means

Expand further: throughput improves as intended, customers retain the expected outcome and realised effort supports the economics. Repeat or enlarge the capacity investment where the next constraint appears.

Adapt: throughput improves but rework rises, or the hire remains blocked by founder decisions. Repair the workflow or decision rights before repeating the hiring pattern.

Pause further additions: demand weakens, performance crosses the agreed quality boundary or the financing assumptions fail. Reassess the next commitment while meeting existing obligations.

The founder decision: invest now in the demonstrated constraint. Keep monitoring the uncertain parts. Evidence-led entrepreneurship includes acting decisively when delay is the weaker choice.

Your one-page decision record

Use this for a consequential choice. Replace the prompts with your own answers. Keep the record short enough to revisit; link to detailed evidence where needed. Do not calculate a weighted readiness score.

Field	Complete before committing
Decision and date	We are deciding whether to ... by The decision owner is
Commitment	Money, people, founder attention and obligations: What is hard to reverse:
Alternative and delay	Best alternative use of these resources: Cost of waiting:
Current constraint	The outcome that is blocked: The limiting factor: Relevant principle:
Critical claim	This step only makes sense if Strongest alternative explanation:
Evidence	Observed: Inferred: Unknown: Strongest evidence against our preference:
Next investigation	We will learn ... through Owner: Maximum cost: Decision date:
Action branches	Proceed if Change the plan if Stop this route if An inconclusive result means
Decision and review	We chose ... because We expect We will review on ... or earlier if

Before approving the commitment

Check the consequences for customers, delivery, economics and cash. Check whether a critical right, person, supplier or partner can prevent completion. Ask whether the evidence is strong enough for this specific commitment and whether delay would improve or worsen the choice.

At the review

Compare what happened with what you expected. Separate a poor decision from an unfortunate outcome and a good decision from luck. Record what changes in your model of the business and who owns the next step.

The output is a decision and a learning loop. A completed form without an owner, action or review trigger adds little value.

What supports this guide—and what does not

This guide combines empirical research, strategic theory and practitioner experience. Its nine principles and operating templates are a synthesis for founders. They have not been tested together as an intervention, and there is no defensible general percentage by which they increase the chance of an exit.

The strongest empirical foundations

Disciplined entrepreneurial learning can be taught. Camuffo et al. studied 759 firms in four randomised trials. Their approach affected idea termination and the focus of strategic changes. This supports structured reasoning under uncertainty; it does not establish a fixed improvement in IPO probability. [\[1\]](#)

Scaling timing deserves scrutiny. Lee and Kim studied more than 38,000 US startups. Early scaling was associated with higher failure without an observed compensating benefit in successful exits. The study uses job postings as a proxy and is observational: it does not prove that early hiring itself causes failure. [\[2\]](#)

People and operating capability matter. Research on the premature death of early employees finds persistent firm-level effects. A separate randomised management intervention in established Indian textile businesses improved performance. These settings support attention to organisational capability, while limiting what can be inferred for a particular young startup. [\[5\]\[7\]](#)

How to interpret the sources

Randomised studies support causal statements about a specific intervention and population. Natural experiments depend on identification assumptions. Observational studies can reveal useful patterns without eliminating selection or alternative explanations.

Strategic theory explains mechanisms such as focus, complementary assets and value capture. Practitioner essays supply useful methods and examples, but can overrepresent winners. An activity seen in successful companies may be a consequence of their success rather than its cause.

The business examples, action branches and decision record are practical applications of these ideas. They are not research-derived universal thresholds. Apply them to your customer, technology, development cycle, capital needs and consequences of error.

What success means here

Success means building enduring customer and business value while reaching an outcome consistent with the founders' ambitions and the capital they accept. Survival, profitability, rapid growth, a sale and an IPO are different outcomes. Market changes, timing and chance continue to matter.

The revision used simulated founder and venture-developer perspectives spanning industrial and rapidly scaling businesses. These were editorial tests, not external interviews. They sharpened the focus on practical decisions, industrial realities and decisive action when evidence supports growth.

Sources and strength of evidence

The numbered citations in the text are clickable. The notes below explain the type of support each source provides. Sources consulted on 23 September 2026.

1 Camuffo et al. (2024)

A scientific approach to entrepreneurial decision-making: Large-scale replication and extension

Strategic Management Journal 45, 1209–1237. Four randomised studies involving 759 firms. Supports more focused search and idea termination; does not establish a fixed improvement in the probability of an IPO.

2 Lee and Kim (2024)

When do startups scale? Large-scale evidence from job postings

Strategic Management Journal 45, 1633–1669. More than 38,000 US startups; job postings serve as a proxy for the start of scaling. Observational research.

3 Koning, Hasan and Chatterji (2021)

Experimentation and startup performance: Evidence from A/B testing

Authors' version consulted, dated 10 September 2021. Longitudinal research on digital startups and A/B testing. The study setting limits the transferability of its findings.

4 Azoulay et al. (2020)

Age and High-Growth Entrepreneurship

American Economic Review: Insights 2, 65–82. US administrative data on founders and companies; age and industry experience are predictors, rather than randomised interventions.

5 Choi et al. (2023)

Early Joiners and Startup Performance

Federal Reserve FEDS 2023-012, working paper consulted. A natural experiment involving premature deaths of early employees; not a direct test of a recruitment method.

6 Edmondson (1999)

Psychological Safety and Learning Behavior in Work Teams

Administrative Science Quarterly 44, 350–383. Study of 51 teams in a manufacturing company. Links psychological safety and learning behaviour; not a startup-specific test of success.

7 Bloom et al. (2013)

Does Management Matter? Evidence from India

Quarterly Journal of Economics 128. Randomised management intervention in Indian textile businesses. Causal evidence within this setting, rather than an estimate for young technology companies.

8 Adner and Kapoor (2010)

Value creation in innovation ecosystems

Strategic Management Journal 31, 306–333. Research on nine generations of lithography technology, 1962–2005. Distinguishes dependencies on suppliers from those on complementary innovations.

9 Teece (1986)

Profiting from technological innovation

Research Policy 15, 285–305. Original theoretical article with cases on value capture, imitation and complementary assets.

Sources and practitioner frameworks

10 Porter and Harvard Business School

Strategy Explained

Original institutional explanation of strategic positioning, trade-offs and the fit between activities. A strategic framework, rather than a measurement of effects in startups.

11 Steve Blank (2010)

What's a Startup? First Principles

Practitioner framework describing a startup as a search for a repeatable and scalable business model. Distinguishes searching from executing.

12 Paul Graham (2013)

Do Things That Don't Scale

Practitioner essay on personal customer acquisition, close attention to early users and manual delivery. Experience-based knowledge using selected examples.

13 Paul Graham (2015)

Default Alive or Default Dead?

Practitioner essay on cash position, cost development and dependence on further funding. Not a universal standard for capital-intensive businesses.

14 David Skok

SaaS Metrics 2.0 – Detailed Definitions

Practitioner definitions of acquisition costs, customer lifetime value, gross margin and payback period. Useful for measurement discipline; SaaS benchmarks are not universal laws.

15 Sarasvathy (2006)

The Affordable Loss Principle

University of Virginia Darden teaching note UVA-ENT-0075. A framework for affordable loss and action under uncertainty, without requiring precise predictability.

16 Euronext (2026)

IPO Guide

Primary market source on preparation, the listing process and life after listing. Institutional practice guidance, rather than evidence that an IPO is the best outcome for every founder.

17 US Securities and Exchange Commission

Ready to Go Public

Version consulted updated on 24 April 2026. Preparation for governance, reporting, resources and leadership. US context; not a Dutch legal checklist.

Questions or feedback during the pilot?

Guido van Gageldonk | Venture developer, BOM

[Contact Guido via the BOM AI-tech page](#)