

THE IN A HURRY SERIES

Startups *in a Hurry*

Idea to MVP, fast

Max Brocklesby

BOOKSINAHURRY.COM

Contents

The Whole Thing in One Page 4

Why You Should Care 6

The Core Ideas 8

How It Actually Works 21

What People Get Wrong 30

Use It 35

Terms 40

Go Deeper 43

Notes and Sources 44

Bibliography 49

The Whole Thing in One Page

A startup is easy to romanticise because the visible result is a product. Someone has an idea and launches. The harder reality is that the product is only one answer to a chain of unanswered questions. Who has the problem? How costly is it? Who can buy? What will make them switch? Can the solution work safely and reliably? Can it be delivered for less than customers will pay? Can the company reach enough of them before the cash runs out?

Fast founders arrange those questions so mistakes are discovered while they are still cheap.

Think of a startup as climbing an evidence ladder. The first rung is **problem evidence**: a specific group repeatedly experiences a problem costly enough to change behaviour. Next comes **demand evidence**: people cross some real threshold to get the proposed value, whether by paying, booking, signing, supplying data, committing time or approving a pilot. Then comes **usage evidence**: the product delivers the intended result and people return when the problem returns. **Economic evidence** asks whether price, acquisition, delivery, support and cash timing can form a business. **Scale evidence** asks whether growth remains healthy when the founders can no longer carry the company by hand.

An MVP belongs near the bottom of that ladder. It is not a miniature final product or permission to ship rubbish. It is the smallest credible arrangement that tests the most dangerous uncertainty. Sometimes that is software. Sometimes it is a mock-up, a manual service, a paid pilot, a prototype on a bench, a regulatory conversation, a supply agreement or a pre-order. In medicine, aviation, finance and other high-consequence settings, safety and compliance are part of viability, not luxuries to add after learning.

The founding team is another early experiment, with expensive mistakes. Cofounders need tested working relationships, explicit motives, decision rights, ownership terms and clean intellectual property. A product can pivot.

A broken founder arrangement is harder to reverse.

Customer discovery helps, but conversation is weak evidence until it is tied to behaviour. People are generous with hypothetical praise. Ask what they did the last time, what they paid, what they tried, who approved it and what stopped them. Then increase the cost of saying yes. A click is stronger than a compliment. A booked call is stronger than a click. Repeated paid use is stronger than both.

Product-market fit is what happens when those forms of evidence begin to reinforce one another. A defined customer gets value, returns, pays, recommends, renews or expands, and can be served without the economics collapsing. There is no universal score that proves fit across a consumer app, a hospital device and a construction marketplace.

Capital, hiring and growth follow the same logic. Money buys time and capability, but also obligations. Every employee, lease, contract, investor promise and public dependency reduces the freedom to change course. Scale amplifies whatever is true. If the model works, growth compounds it. If the model leaks, growth makes the leak larger.

So move quickly where a decision is cheap to reverse and the evidence is informative. Move carefully where the commitment is durable, regulated, expensive or difficult to unwind. The real meaning of fast is not frantic output. It is reaching a better decision before you spend the right to change it.

That is the book.

Why You Should Care

In late 2010 Joel Gascoigne wanted to know whether people would pay to schedule social-media posts. He did not begin by constructing a broad publishing platform. He put up a small landing page explaining the idea. Interest led to a pricing page. Only after he saw enough signal did he build the first narrow version of Buffer. Within days of launch, the product had its first paying customer. Gascoigne later described the sequence as going from idea to paying customers in seven weeks.

The detail matters because the landing page was not magic. It answered one question cheaply: will enough people lean towards this promise to justify the next piece of work? The page could not prove retention, product quality, support cost or a durable market. Each later step had to earn the next one.

Startup stories usually remove that discipline. Retrospectives compress years of uncertainty into a clean line from insight to success. The first product looks inevitable, funding appears to certify the idea, and the mess is edited out because uncertainty makes bad mythology.

For a founder, the edit is dangerous. Most early decisions have asymmetric costs. Choosing a name badly is irritating. Hiring ten people against a false assumption can be fatal. Building a prototype nobody wants wastes weeks. Signing a five-year lease, entering a regulated market without a viable approval path, or raising capital whose expected outcome is larger than the market can support can trap the company inside its own earlier optimism.

The logic matters outside venture-backed technology. A restaurant testing a second site, a manufacturer launching a new line and a scientist commercialising a discovery face different risks, but the common problem is the same: an important fact is unknown and the next commitment can reveal it or bury it.

The public image of startup finance is also skewed. Venture capital is important for a narrow class of businesses that can use large amounts of

risky capital to pursue exceptionally large outcomes. It is not the default funding source for new firms. Research on young US firms has shown substantial use of debt and founder resources, while UK guidance now presents equity, debt, grants and revenue as different tools rather than a single ladder towards venture funding. A business can be excellent and unsuitable for VC. A company can raise VC and still be a bad business.

Nor does failure have the clean meaning startup culture gives it. Official UK business demography shows that 38.4 per cent of businesses born in 2019 were still active five years later. That is a sobering number, but it is not a venture-startup failure rate, and closure is not always catastrophe. Some firms are sold, merged, voluntarily wound down or replaced by better opportunities. Conversely, a company can remain legally alive while destroying time and capital. The useful question is not whether failure is glorious or shameful. It is whether the next pound and the next month are buying evidence strong enough to justify themselves.

This book therefore covers the entire early venture sequence promised by the title: idea selection, cofounders, discovery, MVP design, first users, product-market fit, business model, funding, hiring, scaling, pivots and stopping. The neighbouring books on Marketing, Sales, Negotiation and Management go deeper into those specialist functions. Here they appear only where a founder must decide what to prove next.

The aim is practical without pretending there is one recipe. A consumer software product can test demand in hours. A medical device may spend years proving safety, performance and regulatory acceptability before broad commercial demand can be tested. A marketplace must persuade two sides. Hardware may need tooling and working capital before scale. In each case, the question remains useful: what evidence must exist before this commitment becomes rational?

Answer that well and speed becomes less mysterious. You stop trying to look like a startup and start making the unknowns smaller.

The Core Ideas

1. Start With a Problem That Already Costs Something

Founders usually begin with a solution because solutions are mentally satisfying. A scheduling app, a new battery, a marketplace for tradespeople, a diagnostic tool. A solution can be imagined in an afternoon. A market cannot. It contains people with habits, budgets, competing priorities, gatekeepers and reasons to leave an imperfect system alone.

The first serious task is therefore to find evidence that the problem has a cost before your product exists. Cost is broader than money. It can appear as hours of labour, lost sales, dangerous delay, uncertainty, compliance exposure, embarrassment, rework or an opportunity repeatedly missed. The strongest early clue is often a workaround. People maintain absurd spreadsheets, employ staff to copy data between systems, build internal tools, pay consultants, tolerate queues or break their own process because the alternative is worse. Workarounds are valuable because behaviour has already crossed a threshold. The customer is spending something now.

Severity alone does not create a market. A painful problem can still be commercially useless. The person suffering may not control a budget. The benefit may be diffuse while the purchase cost lands on one department. The buyer may face switching risk larger than the promised gain. Regulation, procurement or integration may make adoption so slow that the startup runs out of cash before the customer can say yes.

In business markets, map the whole decision chain. The user, buyer, approver, technical gatekeeper and person carrying the downside can all be different. A clinician may want the tool, an IT team may reject it, procurement may demand a tender, and a hospital executive may carry the legal risk. Interviewing the happiest user can produce a completely false picture of the sale.

Narrowness helps because patterns are easier to see when the customers

resemble one another in the ways that matter. “Small businesses” is a population category. “Independent veterinary practices with three to ten clinicians that manually chase overdue insurance claims” is a starting segment. The first description hides buying behaviour. The second gives you a place to look for repeated pain, a route to the customer and a plausible definition of success.

Then ask why the opportunity exists now. Startups are often made possible by a changed constraint: a component becomes cheap, an API appears, regulation changes, a new behaviour becomes normal, a distribution channel opens, or incumbents become too slow for a new customer expectation. Timing does not rescue a weak problem, but it explains why a problem that existed for years may suddenly support a new venture.

Founder knowledge matters for the same reason. Domain experience can provide access, language, technical understanding and an instinct for where the expensive friction hides. It can also create blindness. An insider knows why everything is done the current way and may mistake historical accidents for laws of nature. An outsider may see the absurdity but underestimate the constraints. The useful founder advantage is neither ignorance nor experience by itself. It is access to a better hypothesis and the ability to test it.

A good first sentence therefore names a customer, a recurring problem and the costly consequence. It does not need to contain the final product.

Market size comes after this specificity, not before it. A top-down claim that a sector is worth £20 billion tells you almost nothing about whether a new entrant can reach a viable first market. Build from the bottom up: count the customers resembling the segment you have chosen, estimate a plausible price and frequency, then ask what share could be reached through channels that exist. The number will be rough, but its assumptions can be inspected. “One per cent of a huge market” is not a strategy because it explains neither why the one per cent chooses you nor how you find them.

The best early markets are often small enough to understand and large

enough to matter. They give the startup a dense learning environment. Expansion can come later through adjacent customers with a related problem. “We help independent physiotherapists reduce the time spent producing insurer reports” leaves room to discover whether the answer is software, a service, a template or something else. “We are building an AI platform for healthcare” begins with an object and asks the market to supply a reason for it.

The difference is more than wording. One sentence points towards evidence. The other points towards construction.

2. Turn the Idea Into a Stack of Assumptions

Every startup plan contains hidden sentences beginning with “we assume”. We assume this customer feels the problem. We assume they can switch. We assume the buyer can approve the price. We assume the technology reaches the required accuracy. We assume a supplier can deliver. We assume regulation permits the workflow. We assume enough customers can be reached at a cost the margin can bear.

Writing those assumptions down changes the venture because it separates what has been observed from what is merely coherent. A spreadsheet can make a five-year forecast look exact while every input remains invented. Early precision is often cosmetic.

The useful question is not which task comes next. It is which uncertainty deserves to be attacked next. Rank assumptions by two dimensions: how uncertain they are and how damaging it would be if they were false. A beautifully designed checkout flow deserves little attention if the business dies when a regulator rejects the underlying service. A perfect technical prototype is premature if the only likely buyers have no budget. A large demand campaign is wasteful if the product cannot achieve a safety threshold that no amount of enthusiasm can soften.

This is where the evidence ladder begins. **Problem evidence** establishes that a defined group already pays a cost. **Demand evidence** shows that people will cross a real threshold for the proposed value. **Usage evidence**

shows that the offer delivers value and survives contact with routine.

Economic evidence shows that acquisition, price, fulfilment, support and cash timing can fit together. **Scale evidence** shows that the system still works when volume rises and founders stop personally rescuing every edge case.

The rungs are not a universal chronology. A deep-technology company may need technical evidence before meaningful demand evidence because no credible buyer can evaluate a physically impossible promise. A regulated venture may need an early legal or regulatory path. A marketplace often needs supply and demand evidence in parallel. The ladder is a hierarchy of confidence, not a bureaucratic stage gate.

The important move is to prevent weak evidence from impersonating strong evidence. Ten enthusiastic interviews are not ten customers. A waiting list is not retention. Revenue from bespoke consulting is not proof that software will scale. A term sheet is not product-market fit. A rapidly growing user count bought through subsidies is not an economic model.

Research supports the value of disciplined hypothesis testing, with limits. In a randomised trial involving 116 Italian startups, Arnaldo Camuffo and colleagues taught one group to formulate predictions and test hypotheses more rigorously. The treated entrepreneurs were more likely to make informed pivots and performed better on the study's early measures. That does not prove founders should imitate laboratory science mechanically. Markets are strategic, samples are noisy and important opportunities may be hard to express as clean experiments. It does show that explicit theories and tests can improve decisions relative to undirected intuition.

A test needs a decision rule before the result arrives. What would make you continue, change the hypothesis or stop? Without that, every outcome can be narrated as encouraging. Founders are excellent lawyers for their own ideas. They explain away weak conversion, redefine the customer, blame the copy, change the metric and call persistence resilience.

Good evidence narrows the number of stories you can tell yourself.

A useful test portfolio also distinguishes **learning speed** from **decision speed**. You may run several cheap tests in parallel, but you should not let each result trigger an immediate strategic lurch. Noisy markets produce false positives. Combine evidence where the claims are related, repeat a result when the cost of being wrong is high, and reserve irreversible decisions for patterns rather than anecdotes. Fast companies often look calm from the outside because they compress learning without changing direction every afternoon.

3. Founders Are the First Irreversible Product Decision

A prototype can be thrown away. A founder relationship leaves ownership, control, intellectual property and emotional residue even after it breaks. That makes the founding team one of the earliest commitments that deserves slower thinking.

The romantic test is whether two people share a vision. The practical test is whether they can work through ambiguity when the work becomes repetitive, status becomes unequal and one person's preferred answer loses. Compatibility is easier to claim than to observe. Work together before you incorporate where possible. Build something, sell something or run a difficult project against a real deadline. You want evidence about pace, standards, candour, reliability and conflict.

Complementary skills help, but the cartoon pairing of “business founder” and “technical founder” is too crude. Early companies need somebody to understand the problem, somebody to make the solution real, somebody to bring it to customers and enough shared judgement to decide what not to do. Those capabilities can sit in one person, two people or a wider founding group. The weakness is a missing function nobody owns.

Motives need the same treatment. One founder may want a profitable independent company. Another may want the fastest route to a billion-pound outcome. One may treat control as precious; another is willing to exchange it for capital and speed. Noam Wasserman's work on founder dilemmas

made the trade between control and financial upside famous, but the larger lesson is that unstated motives eventually become governance disputes.

So make the awkward conversations early. What does full-time mean? How long can each founder go without salary? Who decides product, hiring and fundraising? Which decisions require unanimity? What happens if one person wants to sell and another does not? How will equity vest? What happens if someone leaves? Who owns code, designs, patents, data and other intellectual property created before and after formation?

There is no universal correct equity split. Equal ownership can reflect equal commitment and simplify trust. Unequal ownership can reflect different risk, contribution or timing. What matters is that the split comes from a considered view of the future rather than a desire to avoid discomfort today. Vesting is common because it reduces the damage of granting a permanent claim to someone who departs early. The specific legal structure varies by jurisdiction and deserves proper advice.

Solo founders remove one category of conflict and concentrate every blind spot. They can move with unusual coherence but need deliberate sources of challenge and enough capability to cover the first critical functions. Cofounders add capacity and information, plus another human system that must be managed. Neither structure deserves mythology.

Early employees matter sooner than many founders expect. Research using US administrative data has found persistent negative effects when startups unexpectedly lose founders or employees who joined in the first year, evidence that important organisational knowledge sits in people before it sits in process. The lesson is not to hire a large team immediately. It is to treat the first few people as part of the venture's design rather than interchangeable labour.

A startup asks customers to trust a product that is still changing. It cannot afford for the people changing it to discover, under pressure, that they never agreed what company they were building.

4. Customer Discovery Is an Evidence Problem

“Talk to customers” is good advice and dangerously incomplete. A badly designed conversation can strengthen the wrong belief because people are polite, hypothetical and eager to be helpful.

Do not begin by pitching the solution and asking whether they like it. Begin with a recent event. “Tell me about the last time this happened.” Reconstruct the sequence. What triggered the problem? What did the person do first? Which tools were involved? Who else entered the process? What did the delay cost? What was tried before? What was paid for? What prevented a better fix?

Past behaviour is not perfect evidence of future adoption, but it is harder to manufacture than hypothetical enthusiasm. Artefacts can be better still. With permission and appropriate care, look at the spreadsheet, handover note, invoice, form, queue, screenshot or calendar that carries the work. People omit steps they have stopped noticing. The artefact preserves them.

Sampling matters. Friends and founder networks are convenient and often unusually forgiving. Online communities can over-represent enthusiasts or complainers. A corporate innovation team may hear from users who love the problem but have no power to buy. Seek people who bought alternatives, abandoned alternatives, refused to buy, control budgets and can block adoption. Negative cases are information.

The interview should end with a stronger test, not a longer transcript. Ask for an introduction to the buyer. Request access to representative data. Book a follow-up. Offer a paid diagnostic. Ask for a deposit, pre-order or signed pilot where those are appropriate and lawful. The aim is to increase the cost of agreement gradually.

Think of evidence as having friction. A compliment costs almost nothing. An email address costs little. Time spent configuring a trial costs more. Internal procurement effort costs more again. Payment, renewal and repeated use cost the customer enough that the signal changes character. The exact order varies by business, but the principle is robust: behaviour with

consequences is harder to fake than stated intention.

Be careful with letters of intent and pilots. A non-binding letter can represent little more than goodwill. A pilot may be championed by one employee with no path to procurement. A deposit may be refundable. The meaningful question is what organisational threshold was crossed. Did the customer expose data? Allocate budget? Involve security? Schedule staff? Put reputation behind the decision? Different commitments answer different questions.

Customer discovery can also mislead by making the present too powerful. If you ask people to design the future, they usually describe improved versions of what they already know. Founders still need judgement about technologies, behaviours and combinations customers cannot yet imagine. The purpose of discovery is not to outsource invention. It is to understand the constraints any invention must overcome.

This is where evidence and vision meet. Vision decides which future is worth attempting. Discovery tells you which parts of your story reality has already rejected.

5. An MVP Is the Smallest Credible Test, Not the Smallest Product

The phrase minimum viable product has caused years of avoidable confusion because all three words invite abuse. “Minimum” becomes an excuse for missing essentials. “Viable” becomes whatever the founder can persuade someone to tolerate. “Product” implies that the test must be a reduced version of the final thing.

A better definition is functional: the MVP is the smallest credible arrangement that can change your decision about a major assumption.

Sometimes the right MVP is a landing page. Buffer began that way because the early question was whether enough people cared about scheduled posting to justify building a narrow tool. Sometimes the answer is a concierge service in which founders deliver the value manually. The manual

work exposes the customer's process before software freezes it. A Wizard of Oz test presents a product-like front while humans perform work behind it, useful when the interaction matters more than the automation at first.

A prototype answers a different question. It may test whether a customer can understand a workflow, whether a mechanism performs, whether hardware fits an environment or whether a technical threshold is reachable. A pre-order or paid pilot tests commitment. A supply-side MVP may establish whether enough providers will participate before a marketplace recruits demand. A regulated venture may begin with a narrow feasibility study and a regulatory path rather than a public launch.

The choice should follow the risk. If desirability is unknown, do not spend months automating. If technical feasibility is unknown, a beautiful mock-up is nearly irrelevant. If adoption depends on integration with a hospital record system, a prototype isolated from that constraint may prove the wrong thing. If the danger is trust, a deliberately rough product can poison the test by creating a second reason for rejection.

This is why “embarrassingly early” is poor universal advice. Some imperfections are harmless. Others are the experiment. A consumer scheduling tool can survive an ugly settings page. A payroll service cannot casually miscalculate wages. A medical device cannot treat safety as a later feature. Cybersecurity, data protection, legal permission and physical safety may be minimum requirements for a credible test.

An MVP should therefore have one primary question, a defined customer, a success signal and a time limit. It should also specify what it will not test. A landing page may show message-level demand while saying nothing about retention. A manual service may prove customers value the outcome while hiding whether automation is possible. A prototype may prove technical performance while saying nothing about price.

The common failure is building a bundle of features that produces an ambiguous result. If users do not return, was the problem weak, onboarding confusing, performance poor, price wrong or the chosen customer

unsuitable? The more hypotheses packed into one release, the more stories can be invented afterwards.

A good MVP reduces ambiguity faster than it increases functionality.

6. Product-Market Fit Is Repeated Choice

Early startups are rich in positive signals that do not add up to a business. Press coverage, waitlists, investor meetings, awards, app downloads and compliments can all be real without showing that customers repeatedly choose the product when they are free to stop.

Product-market fit is the name given to the moment, or more often the period, when the evidence starts reinforcing itself. A defined customer reaches the promised value. The problem returns. The customer returns too. Payment, renewal, referral or expansion follows in a pattern that makes sense for the business. The team spends less energy dragging people towards the product and more energy serving demand that has its own momentum.

Retention is often the cleanest lens because it asks whether the product survives after novelty. For a daily tool, useful cohorts may be measured over days or weeks. For annual compliance software, the cycle may take months. A consumer marketplace may care about repeat booking on the natural interval of the underlying need. There is no single retention percentage that travels across these businesses.

Cohorts matter because total growth can hide leakage. Suppose a company acquires a thousand users each month while almost every previous cohort disappears. The chart rises, yet the business is buying replacements for people it failed to keep. A smaller group with stable repeat behaviour may contain much stronger evidence.

Payment strengthens the case but does not settle it. A founder can sell bespoke work that future employees cannot reproduce. A customer can buy once from curiosity and never renew. A subsidised product can grow while losing money on every transaction. Revenue must be connected to the unit

that creates it.

That means asking basic economic questions earlier than founders enjoy. What does one sale contribute after the costs that rise with serving it? How long does cash remain tied up before payment arrives? What does it cost to acquire a customer through a channel that could plausibly continue? How much support does the product consume? What happens when founder labour is priced as labour rather than treated as free?

Practitioner methods can help without becoming laws. Sean Ellis's strongest disappointment response, later used publicly by Superhuman, gives founders a way to identify the users who would most miss the product. The famous 40 per cent threshold is a heuristic from practitioner experience, not a scientific boundary between fit and failure. It becomes useful when combined with behaviour: who are those disappointed users, what value do they get, how do their cohorts behave, and can more people like them be reached economically?

Fit is therefore segment-specific and mechanism-specific. A startup can have strong fit with one narrow customer and weak fit everywhere else. Averaging those groups together destroys information. Early expansion should begin by understanding the people for whom the product already feels unusually good.

The word fit can sound final. It is not. Competitors arrive, channels saturate, regulation changes and customer needs move. Product-market fit is evidence strong enough to justify a different mode of operation, not a certificate of permanent safety.

One useful diagnostic is to ask what breaks if acquisition doubles tomorrow. If the answer is merely server capacity, the company may be ready to add infrastructure. If the answer is onboarding, support, refunds, manual founder intervention and poor retention, growth is likely to multiply unresolved product problems. Fit is partly visible in what demand does to the organisation. Healthy demand creates pressure around capacity. Unhealthy demand creates pressure around persuasion and rescue.

7. Let Evidence Earn Irreversible Commitment

At the beginning, a startup's main structural advantage is freedom. There are few employees, few systems, few customers depending on it and few contracts that cannot be changed. The company can learn because it has not yet built much that learning would force it to destroy.

Growth consumes that freedom. Hiring turns a hypothesis into salaries and careers. A lease turns an address into a fixed cost. Inventory turns a forecast into cash sitting on shelves. Venture funding creates ownership, governance and expectations about the size and speed of the outcome. A long enterprise contract can force the roadmap towards one customer. Public reliability promises can make rapid experimentation impossible in parts of the system.

None of these commitments is bad. A company becomes real by making them. The danger is the order.

Research on scaling gives the mechanism some empirical weight. Saerom Lee and J. Daniel Kim studied more than 38,000 US startups using millions of job postings and treated early hiring into managerial and sales roles as evidence of scaling. Early scaling was associated with greater failure risk in their setting, consistent with a trade-off between moving quickly and committing resources before enough has been learned. The result is not a command to grow slowly. Some markets reward speed, network effects or early capacity. It is evidence that commitment risk is real.

Hiring should therefore answer a recurring constraint. If founders have ten qualified opportunities and repeatedly lose them because implementation capacity is the bottleneck, an implementation hire may be earned. If the founders dislike prospecting and recruit a sales team before anybody has found a repeatable sale, payroll has been used to hide uncertainty.

Funding should be judged the same way. Capital is useful when it buys a milestone that cannot be reached efficiently from revenue or founder resources: a technical programme, inventory, certification, marketplace liquidity, manufacturing capacity or a distribution push after the economics

are understood. It is dangerous when it merely delays a hard answer.

Venture capital deserves special care because its economics shape the company. VCs build portfolios expecting a minority of investments to generate unusually large returns. Research on VC decision-making shows that investors evaluate teams, markets, products and deal structures through the lens of that portfolio. Their willingness to invest is meaningful evidence about investability. It is not proof that customers have validated the business.

Pivots fit the same framework. A pivot is a change to a major hypothesis after evidence damages the current one while leaving enough assets or insight to justify another attempt. Research on entrepreneurial pivots shows that strategic change involves reallocating attention and resources, not merely renaming indecision. State what failed, what was learned and what the new hypothesis predicts. Otherwise movement becomes a way to avoid stopping.

Stopping is a legitimate outcome. Sunk cost makes founders treat past effort as a reason for future spending even when the new pound faces a bad expected return. The decision should look forward. What evidence could still change the conclusion, what would it cost to obtain, and what is the opportunity cost of another six months?

This closes the loop begun with the first idea. Startups can learn quickly because little is fixed. The successful company eventually must fix many things: people, systems, promises, capital and process. Evidence is what should decide the order in which freedom is surrendered.

Scale is not the moment uncertainty ends. It is the moment being wrong becomes expensive faster.

How It Actually Works

Write the first risk map

Begin with one page. Put the customer and costly problem at the top. Underneath, write the assumptions without which the venture fails: the user experiences the problem often enough; the buyer controls budget; the proposed result is possible; the route to market is reachable; the customer will switch; the price can exceed the costs of delivery; any necessary regulation, safety or supply condition can be met.

Do not score the list to three decimal places. Mark what is known, what is inferred and what is unknown. Then circle the assumption that combines high uncertainty with high damage. That is the first candidate for a test.

The exercise often changes the work immediately. A founder planning six weeks of development discovers that the most dangerous question is whether a procurement team will accept the proposed data flow. A hardware team discovers that component cost, not demand, kills the target price. A marketplace discovers that supply will join only after demand exists and demand will join only after supply exists. The risk map turns a vague idea into a set of decisions.

Add one more column: **cheapest credible evidence**. For each major assumption, write the least expensive evidence that would materially change confidence. The answer might be five procurement interviews, a bench test, a paid pilot, a supplier quote or a regulatory opinion. This prevents the team defaulting to product development as the universal response to uncertainty. Code is only one way to learn.

Give the risk map a review date. Assumptions can change as evidence arrives, competitors move or regulations shift. A risk that was secondary in week one can become the main constraint in month three. The map is a living, recurring operational decision tool, not a founding document to admire later.

Find the first narrow market

Describe the initial customer by situation, not aspiration. “Busy professionals” is nearly useless. “Self-employed bookkeepers who reconcile more than twenty client bank accounts each month and chase missing receipts through email” has a workflow, frequency and reachable population.

Look for people already paying the problem tax. Search job descriptions, product reviews, support threads, industry groups, tenders, specialist forums and competitor complaints to learn the vocabulary. This desk work is reconnaissance. It tells you where to ask better questions; it does not replace customer contact.

Recruit outside your immediate circle. Include people who bought an alternative, people who stopped using one and people who rejected it. In organisations, speak to different roles. A user can reveal friction; a budget owner reveals priorities; a gatekeeper reveals reasons adoption dies.

Ten thoughtful conversations may teach more than a hundred shallow surveys, but there is no sacred interview count. Continue until the important patterns repeat and the remaining differences can be explained by role, segment or severity. If every conversation seems different, your segment may still be too broad.

Reconstruct behaviour before discussing solutions

Ask about the last occurrence. Let the customer walk through it. What triggered the work? What happened next? Where did information move? Who waited? Which workaround appeared? What was paid? What went wrong? Which attempt to fix it failed and why?

Then look for stakes. A two-hour annoyance once a year may produce warm interviews and no market. A thirty-minute annoyance every day may support a strong product. A rare event can still be valuable if the consequence is large enough, as in compliance, security or industrial

downtime.

Separate the user from the transaction. Who signs? Who pays? Who can block? What budget category contains the spend? What happens if the buyer does nothing? A problem that cannot survive those questions is not yet a commercial problem.

Put a price on the smallest promise

Price is useful before the product is finished because it forces the value proposition to become specific. You do not need the final price. You need a plausible exchange that exposes whether the buyer treats the problem as economically meaningful.

For a service, offer a narrow paid engagement. For software, test a paid pilot or early plan where appropriate. For hardware, a refundable reservation may reveal more than a mailing list, though it remains weaker than a completed purchase. For enterprise sales, a serious pilot agreement that allocates staff, data and budget can be stronger than an enthusiastic letter from an individual champion.

Avoid fake scarcity and misleading pre-sales. The customer should understand what exists, what does not and what happens to any money committed. An experiment that depends on deception may teach you that deception works rather than that the business does.

Choose the MVP that matches the risk

Now decide what must exist to answer the next major question.

If the risk is **message and demand**, a landing page, sales call or concept demonstration may be enough. If the risk is **workflow value**, deliver the service manually for a few customers and record where the labour goes. If the risk is **usability**, build a realistic prototype. If the risk is **technical feasibility**, isolate the hardest technical claim and test it directly. If the risk is **marketplace liquidity**, recruit one side in a tightly constrained geography or category before widening. If the risk is **regulatory feasibility**, obtain the relevant specialist guidance before designing around assumptions that may

be unlawful.

Keep the MVP narrow enough that you can interpret failure. State the primary hypothesis, the customer, the behaviour expected, the time window and the next decision. A product with twelve new features and one conversion metric is a poor experiment because the cause of the result remains hidden.

Get the first users manually

The first customers are research partners with money attached. Before contacting them, write the smallest sales script that can fail cleanly: who the product is for, the problem it addresses, the concrete result and the next commitment you are asking for. Avoid a ten-minute founder autobiography. Early sales conversations are another form of discovery, and every objection should be classified. “Too expensive” may mean weak value, wrong buyer, no budget, poor timing or a genuine price constraint.

Keep a simple pipeline even when there are only twenty prospects. Record source, segment, stage, objection, next action and outcome. The purpose is not corporate ceremony. Memory becomes selective quickly, especially when founders are desperate for positive signal. A pipeline reveals whether the same type of customer moves forward and where the sale repeatedly stops.

Recruit the first users through direct channels where you can see who responds and why: personal outreach, communities, industry events, existing professional networks, targeted cold contact or partnerships. Mass acquisition is usually less useful while the team is still trying to identify the right customer and message.

Do work that will not scale. Onboard people personally. Watch them use the product. Fix individual problems. If the service is manual behind the scenes, record each intervention. The goal is to discover which parts of the promised value are repeatable and which exist only because a founder is quietly doing heroic work.

Do not confuse generosity with a permanent operating model. Founder-led support can keep a weak product alive. Mark every manual rescue. If the same intervention appears repeatedly, it may reveal a missing product feature, a customer unsuitable for the segment or a cost the business model must eventually carry.

The first-user process should also create a loss log. Record every person who considered the offer and did not proceed, every trial that stalled, every customer who required a discount and every user who reached value only after founder intervention. Review losses by segment and stage, because ten different objections may hide one upstream problem such as poor targeting or an unclear promise. Early teams naturally study the people who stayed because those customers are available to talk. The people who disappear often contain the sharper information. A simple loss log prevents survivor bias from turning the first happy users into a fictional market.

Instrument the moment value appears

Choose an activation event tied to the promised result. Account creation is rarely enough. A scheduling product might count the first successfully scheduled post. A marketplace might count a completed transaction. A workflow tool might count the first report produced and accepted. The metric should correspond to value, not administrative movement.

Then build cohorts. Group customers by when they started, which segment they belong to or how they entered. Track whether they return on the natural cycle of the problem. This is more informative than a single total-user graph because it shows whether later customers behave better after product changes.

Qualitative evidence still matters. Ask retained customers what they would do if the product disappeared. Ask churned customers what happened. Compare the stories with behaviour. A user who says price was the reason for leaving but barely used the product may be describing the polite reason rather than the causal one.

Do not instrument everything because measurement has a cost. A startup

with fifty dashboards can remain confused if none corresponds to a decision. Keep a small operating set: one measure of reaching value, one of returning, one of paying or commercial commitment, and one of the current bottleneck. Add diagnostic metrics when a primary measure moves. The question behind every chart should be: what decision would change if this number changed? If there is no answer, the metric is probably decoration.

Put the business model beside the product

As soon as real transactions appear, sketch the unit economics. What does the customer pay? Which costs rise when another customer is served? What sales or acquisition effort was required? How long before cash is collected? What refunds, returns, support or implementation work appear after the sale?

Do not force a sophisticated lifetime-value model onto five weeks of data. Early estimates should expose what you do not know. If retention history is short, model several scenarios. If acquisition came entirely from the founders' networks, do not pretend the observed acquisition cost will survive scale.

Cash timing deserves separate attention. A business can have positive margins and still run out of cash because inventory, payroll or customer payment terms pull money forward. Growth can worsen that problem. A physical-product company that pays suppliers sixty days before customers pay can require more cash precisely when sales accelerate.

Decide how the next milestone should be financed

Start with the milestone, then choose the money. What evidence or capability would materially change the venture's risk? How long and how much cash will it take? What financing source fits the asset and the upside?

Revenue is attractive because it validates customers and preserves ownership, but it may be too slow for research, certification, inventory or network effects. Founder money is flexible but limited and personally risky. Debt preserves equity but introduces repayment obligations and may be

unavailable or unsuitable before predictable cash flow. Grants can finance qualifying work without dilution but come with eligibility, timing and administrative constraints. Angel and venture equity can support large uncertain bets but exchange ownership and often governance for capital.

For a venture investor, ask whether the market and outcome can plausibly support the return structure of the fund. A healthy company capable of becoming worth £20 million may be an excellent founder-owned business and a poor fit for a fund that needs rare outcomes hundreds of millions or billions in size. Financing strategy should not distort the company into a category it cannot economically satisfy.

Runway should be calculated from the decision calendar, not from wishful averages. If a regulatory milestone takes nine months and fundraising routinely consumes several more, six months of cash is not six months of useful experimentation. Build a downside case in which revenue is late, hiring takes longer and the next round does not arrive on schedule. The purpose is not pessimism. It is to preserve enough time to make a rational decision before the bank balance makes it for you.

Hire only against repeated work or a decisive experiment

Write down what is preventing progress. If the constraint is founder time spent on a repeated, validated task, a hire may be justified. If the constraint is an unanswered strategic question, hiring a department often makes the question more expensive without answering it.

Early roles should have unusually high information density. They sit close to customers, product and operations, and what they learn should alter the system. Hire people who can own outcomes under incomplete structure, then add specialisation as work becomes repeatable.

A first sales hire is especially revealing. If the founders cannot describe who buys, why they buy and how a sale progresses, they are asking an employee to discover the sales model while carrying a quota. That can work with the right person, but it is a different job from scaling a known process.

Read the evidence as a portfolio, not a trophy cabinet

At regular intervals, lay the evidence out by rung. Problem: how often and how costly is the pain? Demand: what costly commitments have customers made? Usage: do they reach value and return? Economics: does serving them create contribution and manageable cash demands? Scale: which bottleneck appears when volume rises?

Put dates beside the evidence. A pilot agreed twelve months ago under a former budget owner may no longer support the same conclusion. Supplier quotes expire. Regulation changes. Competitors alter customer expectations. Startups can become attached to old proof because obtaining it was difficult. Evidence has a half-life, especially where the market is moving quickly.

Do not add weak signals together to manufacture confidence. A famous adviser, ten thousand free sign-ups and investor interest do not compensate for zero retention. Equally, do not demand mature evidence too early. A two-week-old enterprise product cannot show annual renewal. Judge the evidence that the model could reasonably have produced by now.

Then choose one of four moves. **Continue** when the evidence supports the current hypothesis and the next uncertainty is clear. **Repair** when the problem and customer remain sound but execution failed. **Pivot** when a major hypothesis fails and a credible alternative emerges from what was learned. **Stop** when the remaining evidence would be too expensive to obtain relative to the plausible upside or when a core constraint has no credible route through it.

Earn the right to scale

Before increasing acquisition, run a pre-mortem on scale. If volume triples, which part of the system fails first? Support? Supply? Reliability? Cash? Regulation? Founder approvals? Write the failure modes down and decide which need capacity now and which can be handled manually for another cycle. Premature process is wasteful, but foreseeable collapse is not agility.

Then test growth in increments. Add one channel, one geography or one customer class at a time where possible. A startup that expands product, market, channel and team simultaneously loses the ability to explain what caused the result. Controlled expansion preserves learning while exposing whether the original fit transfers. New segments should have their own retention and economics rather than being blended into the old average.

Scale when demand is repeatable enough that the company knows what it is amplifying. That does not require perfection. It requires a coherent story across customer, product and economics.

For some businesses, the signal is strong retention and referrals. For enterprise software, it may be repeatable sales, successful deployment and renewal. For hardware, reliable manufacture and channel economics matter. For marketplaces, liquidity and repeat transactions on both sides matter. For regulated products, approval and safe real-world operation may be prerequisites before distribution can expand.

Watch the bottleneck move. Before fit, the bottleneck is often finding anyone who cares. After fit, it may become implementation capacity, inventory, customer support, reliability, working capital or hiring. Scale should attack the new bottleneck, not copy the organisation chart of a larger company.

The transition from startup to company is gradual. Forecasts become more useful. Roles narrow. Processes become necessary because repeated work should not be rediscovered each week. That is where Management, Marketing and Sales begin to own more of the problem. The startup has not eliminated uncertainty. It has accumulated enough evidence that execution deserves a larger share of attention.

How we know

Entrepreneurship research sees an unusually selected world. Official business registers cover firms that become formal enough to appear in administrative data, while venture databases over-represent companies seeking external finance. Case studies can reveal mechanisms but are poor guides to prevalence, and successful founders are far more likely to be interviewed about methods than founders whose identical methods failed.

That is why this book avoids universal claims such as “90 per cent of startups fail” and treats famous MVP stories as illustrations rather than causal proof. The stronger parts of the model rest on convergence: economic work on experimentation, randomised evidence on hypothesis testing, research on founder and early-employee effects, studies of startup scaling, official survival statistics and practitioner methods tested across many companies. The evidence is strongest for the value of learning before commitment, weaker for any single prescribed sequence or metric. Different technologies, regulations and capital requirements change which test should come first.

What People Get Wrong

“The idea is the valuable part”

A good idea matters, especially where it contains difficult technical knowledge, unusual access or protectable intellectual property. The mistake is treating conception as though it contains most of the venture's value.

Early ideas are bundles of guesses. The customer may differ from the imagined one, the workflow may be wrong, the route to market may be blocked and the feature that seemed central may prove irrelevant. What compounds value is the sequence of discoveries that turns those guesses into a working system.

This is why secrecy can become expensive. Some technologies require confidentiality and proper IP strategy, but many founders hide ordinary ideas

from the people who could disprove them. Fear of copying delays the more immediate threat: building something nobody needs. A defensible venture is usually more than a sentence. It is knowledge, product, distribution, relationships, data, execution and timing accumulated together.

There are exceptions. A patentable invention, proprietary dataset or trade secret can carry substantial standalone value, and premature disclosure can damage protection. The correction is therefore not “share everything”. It is to protect what is properly protectable without turning ordinary commercial uncertainty into secrecy theatre.

“An MVP should be cheap and ugly”

Cheap is useful only if the test remains credible. Ugly is acceptable only when appearance is not part of the thing being tested.

If customers reject a financial tool because it feels untrustworthy, you may not know whether they rejected the underlying value or the rough presentation. If a prototype crashes during the one workflow meant to show benefit, the test creates an alternative explanation for failure. If a medical or safety-critical product is unreliable, the experiment may be unethical as well as uninformative.

Minimum means remove what does not need to exist for this decision. It does not mean remove whatever is difficult. A narrow product can be polished where credibility matters and crude everywhere else. The correct minimum is defined by the hypothesis, not by founder impatience.

The best MVP can therefore cost more than a sloppy version. A reliable bench prototype may be expensive but still minimal if it isolates the technical question that decides the venture. Minimum refers to scope, not necessarily to cash.

“Customer interviews tell you what to build”

Customers are excellent witnesses to their own work and unreliable architects of products they have never seen.

Ask them to reconstruct behaviour, costs, failed attempts, constraints and decision processes. Those answers can reveal what the solution must accomplish. Feature requests are different. A request for “an export button” may conceal a deeper need to move data into an old approval process. Building every requested feature turns a product into an archive of local workarounds.

Discovery therefore informs judgement rather than replacing it. The founder still has to infer which needs are shared, which constraints can be removed and what new behaviour the product can make possible. Listening is not obedience.

The same distinction protects founders from the loudest customer. One large prospect can request enough custom work to bend the entire product around its process. That may be rational if the company has chosen a services model. It is dangerous if the startup believes it is proving a repeatable product while quietly becoming a bespoke contractor.

“Product-market fit is one magic number”

There is no universal product-market-fit threshold. The popular 40 per cent disappointment benchmark is a practitioner heuristic, not a natural constant. A retention curve that looks excellent for a messaging app might be meaningless for annual tax software. A revenue target can look strong while each contract is bespoke and unrepeatable.

Fit should appear in the behaviour appropriate to the model: repeated use, renewal, referral, expansion, shortening sales friction, willingness to pay and economics that survive delivery. The signals need to cohere within a defined segment.

The attraction of one number is psychological. It turns uncertainty into permission. The danger is that founders optimise the permission slip instead of understanding why customers stay.

Fit can also exist before founders can measure it neatly. A tiny enterprise cohort may contain only six customers, but all six may renew, expand and

introduce peers. Small samples demand caution, not automatic dismissal. Behaviour, mechanism and economics still need interpretation.

“Funding validates the startup”

Funding validates one thing directly: investors were willing to fund the startup on agreed terms. A large round can even make a weak company look stronger from the outside because cash buys hiring, offices, marketing and time. Those are capabilities, not customer proof. If the product still leaks users or every sale requires founder heroics, more capital can postpone the diagnosis rather than cure it.

That matters. Sophisticated investors may contribute information, networks and a useful external judgement. But their objective is not identical to the customer's. Venture investors construct portfolios and can rationally back companies with substantial failure risk because a few exceptional outcomes can repay the fund. An investment decision therefore does not prove product-market fit, customer value or sound unit economics.

The reverse mistake is treating bootstrapping as moral superiority. Some businesses cannot reach the next decisive milestone without substantial external capital. The right question is what the money buys, what obligations it creates and whether the company's plausible outcome fits the financing model.

Fundraising can even weaken customer learning by becoming a parallel market for persuasion. Founders may spend months optimising a pitch to investors whose approval feels more immediate than a slow enterprise sale. Capital can extend runway, but the fundraising process itself can consume the attention the runway was meant to protect.

“Move fast and break things”

Speed is contextual. A founder should move fast through reversible choices and fast learning. That is different from moving carelessly through every decision.

A landing page can be changed tonight. A bad equity agreement, unsafe

device, damaged reputation, data breach or factory lease may survive long after the experiment. The more irreversible the commitment, the more evidence it should demand.

This distinction is especially important outside consumer software. Hardware, healthcare, finance, aviation and other regulated fields cannot copy a slogan produced by a social-media era and call the result entrepreneurial courage. The useful version is less cinematic: shorten the path from uncertainty to evidence without creating avoidable irreversible harm.

Speed also has coordination costs. A team that changes priorities faster than customers can experience the product may create motion without learning. Move quickly enough to test the hypothesis, then leave the test intact long enough to observe it.

“Nine out of ten startups fail”

The number is repeated because it is memorable and usually undefined. It also flattens businesses with radically different risk. A bootstrapped consultancy, a drug-development company, a restaurant and a venture-backed consumer network do not share one failure process. Aggregating them creates drama without much decision value. What counts as a startup? What counts as failure? Over what period? Which country, sector and cohort?

Official UK data show that 38.4 per cent of businesses born in 2019 were still active five years later. That does not mean 61.6 per cent “failed” in the venture-capital sense. The dataset covers registered businesses broadly, and business death can include several forms of closure. US establishment survival data use another definition again.

Post-mortem lists claiming one dominant failure cause have the same problem. Startups often fail through interacting mechanisms: weak demand produces poor growth; poor growth drives fundraising pressure; pressure worsens team conflict; the team cuts the wrong work; cash disappears. A single label assigned afterwards can hide the chain. Treat failure statistics

as context, not prophecy.

The practical implication is to build your own failure tree. What would cause demand to disappear? What would cause cash to run out? What would cause founders to split? What technical, legal or supply constraint could stop delivery? These mechanisms can be monitored. A generic failure percentage cannot.

The useful mindset is neither optimism nor pessimism. It is disciplined specificity, tied to mechanisms the team can observe. Name the mechanism that could kill the company early enough that somebody can still test or change it.

Use It

Ask what must be true

When an idea feels exciting, convert it into conditions. What must be true about the customer, technology, price, channel, regulation, supply and timing for the idea to work?

This is useful far beyond startups. A proposed internal project, acquisition, new product line or career move often arrives as a story with the assumptions hidden inside it. Writing the assumptions down exposes which parts are evidence and which parts are hope.

Then rank them. The next action should reduce an uncertainty capable of changing the decision, not merely produce visible activity. This is a better definition of progress than a longer task list.

A useful meeting discipline follows. Instead of asking “what did we ship?”, ask “which belief changed this week?” Shipping matters, but its purpose before fit is partly informational. Teams that cannot name a changed belief may be producing output faster than understanding.

Demand stronger evidence as commitment rises

Match the proof to the size of the irreversible bet. A low-cost prototype can be justified by weak early signal. A full-time hire deserves stronger evidence. A factory, major financing round or regulated rollout deserves stronger evidence again.

This does not eliminate risk. It changes the order in which risk is taken. The principle is close to ordinary prudence: do not spend permanent money on temporary confidence.

You can use the ladder explicitly. What evidence do we have of the problem? What evidence of demand? What evidence of repeated value? What evidence of economics? What evidence that scale preserves the mechanism? If a later rung looks strong while an earlier one is missing, investigate the mismatch.

For example, strong revenue with weak retention may indicate project work, one-off demand or heavy sales effort. Strong retention with weak acquisition may indicate a good product trapped behind an ineffective channel. The mismatch tells you where the next experiment belongs.

Separate compliments from commitments

People want to be kind. Markets have fewer manners. Before any experiment, write the strongest commitment the customer could reasonably make at this stage. For a concept interview it may be access to records or a follow-up with the buyer. For an early service it may be payment. For a regulated enterprise product it may be an approved evaluation plan rather than immediate purchase. This prevents founders from choosing an easy signal after the fact.

Polite interest is abundant; costly commitment is the scarce signal.

When evaluating interest, ask what the other party had to give up. Did they spend money, time, political capital, data, internal effort or reputation? Did they return when the novelty passed? Did they choose you over a real alternative?

This lens is useful in sales, partnerships, hiring and fundraising. “Sounds interesting” belongs in a different category from “I have introduced you to procurement and reserved budget”. Treating the two as equivalent is one of the easiest ways to manufacture momentum.

You can formalise this without pretending the scores are scientific. Write each signal beside the cost borne by the other party. Zero-cost praise sits at the bottom. Time, internal coordination, money and repeat choice sit progressively higher. The exercise forces a team to stop counting every positive response as the same unit.

Look for the bottleneck that moved

A healthy venture changes what constrains it. Draw the customer journey from first contact to repeated value and mark where the largest drop, delay or manual burden sits. Then ask whether fixing that constraint would create another one immediately downstream. A faster sales process can overload onboarding. Better onboarding can expose poor retention. More retention can create support pressure. The movement of the bottleneck is evidence that the system is becoming more real.

Progress is visible when the binding constraint moves. At first the problem may be finding anyone who cares. Then it becomes delivering the value. Later it may be retention, acquisition, implementation, supply, working capital or management capacity.

If the same bottleneck persists while the dashboard grows, the growth may be cosmetic. A company adding thousands of sign-ups while retention remains broken has not moved through the problem. It has made the top of the funnel louder.

Ask each month: what prevented twice as much useful customer value? The answer is more informative than choosing a fashionable metric in advance.

The bottleneck should determine organisational shape too. If technical reliability is constraining growth, more marketing can make the problem worse. If qualified demand is scarce, adding engineering capacity may

produce features for an audience that has not arrived. Resources belong where the system is currently constrained.

Price learning by decision value

The cheapest test is not always the best test. What matters is the cost of obtaining evidence capable of changing what you do.

A £100 landing-page experiment is expensive if it tests message interest when the real uncertainty is a £2 million engineering constraint. A £20,000 technical study can be cheap if it prevents two years of development. A week spent with procurement may be more valuable than a month of user-interface work.

This lens helps prevent false thrift. Founders should conserve cash, but cash exists to buy information and capability. The relevant comparison is the cost of the test against the cost of remaining wrong.

Include time in that cost. Waiting six months for perfect data can be expensive where a market is moving and competitors are learning. Conversely, rushing a test whose signal will be too noisy to guide action wastes both time and money. The best experiment buys enough precision for the decision at hand, not maximum precision. That threshold should rise with the cost of the commitment that follows.

A final check is to ask whether the test itself creates path dependence. A heavily discounted pilot may attract customers who will never accept the intended price. A free marketplace can train users to resist later fees. A custom integration can make one customer look like a market. Experiments should minimise distortion as well as cost, because distorted evidence can be worse than no evidence. When distortion is unavoidable, state it explicitly so the result is not generalised beyond what was tested.

The limits

The evidence ladder can become a fetish. Some opportunities are created rather than discovered. Customers may not be able to express demand for a behaviour that depends on a new technology, social norm or network that does not yet exist. Experiments can favour incremental ideas because incremental ideas are easier to test. Competitors can observe public tests. Small samples are noisy. Founders can choose metrics that flatter their theory. Some decisive evidence arrives only after large commitments have already been made.

The framework also assumes the venture has permission to experiment. Regulation, safety, ethics and contractual duties can restrict what can be tested and with whom. In those settings, the sequence must include specialist review, formal validation or controlled trials before ordinary market learning.

Nor does evidence remove judgement. Two founders can see the same weak signal and rationally make different choices because their opportunity costs, risk tolerance, capabilities and beliefs about the future differ. Entrepreneurship remains decision-making under uncertainty, not a laboratory procedure with an answer key.

The one thing to keep

Do not remember MVP as “build less”. Remember it as **commit less until reality has earned more**.

An idea is a claim about a future customer. A startup turns that claim into progressively harder evidence. People describe a painful problem. Then they act. Then they pay. Then they return. Then the economics hold. Only then does the company deserve more people, more capital, more systems and less freedom to change.

That sequence is why small teams can compete with large organisations. They begin with less, but less is useful when the world has not yet told you what deserves to become permanent.

The founder's job is therefore not to predict the entire company correctly at the start. It is to arrange the next few decisions so that being wrong is cheap enough to survive and being right produces evidence strong enough to act on.

Terms

Activation. The first behaviour showing that a new user has reached the product's intended value, rather than merely registering. The right activation event depends on the product and should predict a meaningful later behaviour where possible.

Angel investor. An individual who invests personal capital in young companies, often earlier and in smaller amounts than institutional venture funds.

Beachhead market. A deliberately narrow first customer segment chosen because its members share a similar problem, buying process and route to reach them.

Bootstrapping. Building primarily from founder resources and operating revenue rather than substantial outside equity. It preserves ownership but can constrain speed, hiring and capital-intensive experiments. It is a financing choice, not a virtue signal.

Burn rate. The net rate at which a company consumes cash over a period. Gross spending and net cash loss are different numbers, so founders should specify which they mean for practical planning purposes.

Cap table. The record of who owns which securities in a company. Founder shares, employee options and investor holdings all change control and future dilution.

Churn. The rate at which customers, users or recurring revenue leave during a period. The denominator and time interval must be defined for the number to mean anything.

Cohort. A group of customers sharing a start date, segment, channel or

other characteristic. Cohort analysis reveals behaviour that aggregate growth can hide, especially whether newer customers retain better than earlier ones.

Concierge MVP. A test in which people manually deliver the promised outcome for a small number of customers, allowing the team to learn the workflow before automating it.

Contribution margin. Revenue from a sale minus costs that rise directly with serving that sale. It is a useful early economic lens but not the same as company profit.

Customer acquisition cost, CAC. The sales and marketing cost required to acquire a customer under a defined method and period. Founder networks and tiny early samples can make CAC look artificially low.

Customer discovery. Structured work to understand customers, problems, alternatives, buying processes and constraints before treating a proposed solution as established fact.

Dilution. Reduction in an existing owner's percentage when new shares or share-like securities are issued. A smaller percentage can still become more valuable if the company grows.

Founder-market fit. An informal description of a founder's relevant insight, access, skill, credibility or persistence in a market. It is a useful question, not a standardised score.

Hypothesis. A claim the venture treats as uncertain and capable of being tested, such as who buys, what they value, what they will pay or whether a technology works.

Letter of intent, LOI. A document expressing an intention to pursue a transaction or commercial relationship. Binding force varies, and a weak LOI can be little stronger than polite interest.

Lifetime value, LTV. An estimate of the economic value generated by a customer over the relationship. Early LTV is fragile because long retention

has not yet been observed.

Minimum viable product, MVP. The smallest credible arrangement capable of producing decision-changing evidence about a major startup assumption. It need not resemble the final product closely.

Product-market fit, PMF. A condition in which a defined market repeatedly chooses and receives value from the product strongly enough that demand, retention and economics begin to reinforce one another.

Pivot. A material change to a major venture hypothesis, such as the customer, problem, channel or business model, made in response to evidence rather than random restlessness.

Pre-order. A commitment to buy before full delivery. Its evidential strength depends on price, refundability, delivery credibility and how much friction the customer accepts. It tests demand, not retention.

Retention. The proportion of a cohort that continues using, buying or renewing over a meaningful interval. Meaningful intervals differ sharply across products.

Runway. The estimated time before available cash is exhausted at the current or planned burn rate. It is a planning estimate, not a guaranteed deadline, and should leave enough time to react before cash is nearly gone safely and irreversibly.

SAFE. A simple agreement for future equity, popularised by Y Combinator, in which investment converts into equity under specified future conditions. Terms and legal treatment require care.

Scale. Committing resources to expand the customer base and the system serving it. Scaling increases upside and commitment risk at the same time, so its timing matters as much as ambition.

Unit economics. Revenue and variable or directly attributable costs analysed at the level of one customer, order, transaction or other useful unit.

Validation. Evidence strong enough to increase confidence in a particular

hypothesis. The word should always prompt the question: validated which claim, with what behaviour, in which customer segment and for how long?

Venture capital, VC. Institutional equity investment in companies expected to pursue unusually large growth outcomes. VC is one financing model, not a synonym for startup finance.

Vesting. A mechanism by which ownership is earned over time or milestones rather than becoming fully unconditional immediately, reducing the damage of early departures and misaligned long-term ownership.

Wizard of Oz MVP. A test that presents an apparently automated customer experience while people perform some work manually behind the scenes, used to test value before building the full system.

Go Deeper

Eric Ries, *The Lean Startup* (2011). Read this for the vocabulary that made build-measure-learn, validated learning and the MVP mainstream. It remains the most accessible starting point for experimental startup practice. Its examples lean towards software and fast iteration, so carry the principles carefully into regulated, hardware or capital-heavy settings. Read it for the experimental mindset rather than for a universal recipe, and notice how easily its vocabulary can be turned into ritual when the underlying hypothesis is forgotten. Pair the famous case studies with less glamorous evidence from your own market.

Steve Blank and Bob Dorf, *The Startup Owner's Manual* (2012). Read this when you want process rather than inspiration. It is long, repetitive and deliberately operational, with detailed treatment of customer discovery and validation. The terminology reflects its era, but the discipline of getting outside the building and testing business-model assumptions remains useful. Use it as a reference manual rather than a book to imitate linearly; the checklists are strongest when adapted to a specific risk. Its greatest value is forcing assumptions about customer, channel and revenue into the open.

Noam Wasserman, *The Founder's Dilemmas* (2012). Read this before treating cofounder, equity and control choices as paperwork. Wasserman uses research and cases to show how early people decisions shape later financing, leadership and ownership. It is especially valuable because these commitments are harder to reverse than an MVP. Read it before incorporation or fundraising if possible, since many of the dilemmas become expensive precisely after the founders believe the awkward decisions are already settled. The cases are US-centred, but the human conflicts travel well across startup systems and financing environments.

William R. Kerr, Ramana Nanda and Matthew Rhodes-Kropf, “Entrepreneurship as Experimentation” (2014). Read this short *Journal of Economic Perspectives* article for the economic model underneath the practical advice. It explains why uncertainty, financing and the option to abandon shape entrepreneurial experimentation. It is academic, but compact and clearer than many startup books about why staged learning matters. It also provides a useful antidote to founder folklore by treating experimentation as an economic process shaped by finance, option value and the ability to abandon weak projects. Read it after the practical books to see the logic underneath their advice. The article is especially useful for understanding why cheap failure can create value when it prevents larger capital from following a weak hypothesis, and why financing affects which experiments are attempted.

Notes and Sources

The word *startup* has no single statutory or scholarly definition. This book uses it operationally for a young venture whose important customer, product, technology, channel or economic assumptions remain unresolved. Steve Blank's influential formulation describes a startup as an organisation searching for a repeatable and scalable business model. The book keeps the distinction useful without pretending that every new bakery executes a known model or every startup must seek hypergrowth.

Sources for the opening and central model

Evidence before commitment. The evidence-ladder model is an editorial synthesis rather than a named framework. It draws on Blank's customer-development tradition, Eric Ries's validated-learning approach, Kerr, Nanda and Rhodes-Kropf's economic account of entrepreneurship as experimentation, and Camuffo and colleagues' experimental work on more scientific entrepreneurial decision-making. The central claim is deliberately narrower than “experiments cause startup success”: explicit hypotheses and staged tests can improve the quality of early decisions when uncertainty is material.

Buffer. Joel Gascoigne's own 2011 and 2020 accounts describe starting Buffer with a landing page to gauge interest, adding a pricing page, building a narrow first product and obtaining an early paying customer. His retrospective describes the path from idea to paying customers as taking seven weeks. The example is used to illustrate sequencing, not to imply that landing pages are universally valid MVPs.

UK survival. Office for National Statistics, *Business Demography, UK: 2024*, reports that 38.4 per cent of businesses born in 2019 survived into 2024. ONS defines survival through continued activity in employment or turnover. This is a broad business-demography statistic, not a venture-backed startup failure rate. The text uses it only to reject unsupported universal claims such as “nine out of ten startups fail”.

Sources for the Core Ideas

Problem discovery and customer behaviour. Blank and Dorf remain the main practitioner source for customer discovery, while Ries provides the broader validated-learning frame. The book's preference for recent behaviour, workarounds and costly commitment is a synthesis of these practices with general evidence principles. It is not a claim that stated preference has no value.

Hypothesis testing. Camuffo, Cordova, Gambardella and Spina ran a randomised controlled trial with 116 Italian startups. The treatment group

was trained to articulate theories, make predictions and test hypotheses more rigorously. Their *Management Science* paper reports better early performance and more informed pivoting in the treated group. The result is bounded by its sample, intervention and early-stage outcomes, so the manuscript does not turn it into a universal scientific method for entrepreneurship.

Founders and early employees. Noam Wasserman's research programme underlies the treatment of founder motives, equity, control and vesting. Choi, Goldschlag, Haltiwanger and Kim use US administrative employer-employee data and premature deaths as a natural experiment. Their final 2025 *Review of Economics and Statistics* article finds persistent negative effects from losing founders and employees who joined in the first year. This supports the importance of early organisational human capital without establishing an ideal founding-team size.

MVPs. Ries popularised the term minimum viable product. Blank and Dorf and Y Combinator's Startup Library provide practical examples of landing-page, concierge, manual and prototype approaches. The present book deliberately reframes MVP around the current uncertainty because the common software-first interpretation travels badly into regulated, safety-critical and capital-intensive ventures. UK Government technology-transfer guidance is consistent with treating intellectual-property, regulatory and commercial constraints as early design inputs rather than later clean-up.

Product-market fit. There is no accepted scientific threshold for product-market fit. The book therefore treats PMF as a segment-specific pattern of repeated choice, value delivery and viable economics. Rahul Vohra's account of Superhuman explains the use of Sean Ellis's disappointment survey and the 40 per cent heuristic. The threshold is identified as practitioner benchmarking rather than a law.

Scaling. Lee and Kim's 2024 *Strategic Management Journal* study uses more than 6.3 million job postings from over 38,000 US startups and identifies early hiring in managerial and sales roles as markers of scaling.

Their results associate early scaling with greater failure risk and interpret the pattern as a trade-off between imitation risk and commitment risk. The book preserves that qualification and does not claim that rapid scaling is always wrong.

Funding. Gompers, Gornall, Kaplan and Strebulaev surveyed 885 institutional venture capitalists at 681 firms about sourcing, selection, valuation, contracting and post-investment activity. Their findings support the separation between investor approval and customer validation. Robb and Robinson's study of new US firms documents substantial use of outside debt and owner-backed financing, countering the public image that startup finance is synonymous with venture capital. British Business Bank guidance currently presents equity, debt, grants and other options as distinct financing tools for smaller UK businesses.

Pivots. Kirtley and O'Mahony's *Strategic Management Journal* work treats pivoting as strategic reorientation involving changes in activities, resources and attention. This supports the manuscript's distinction between an evidence-led pivot and repeated movement without a failed hypothesis.

Sources for the operating sequence

The operating sequence is deliberately constructed for one-hour use. It is not a claim that all ventures move linearly from problem interviews to MVP to PMF. Deep technology may require technical proof first. Healthcare and other regulated fields may require safety, clinical or regulatory work before ordinary market testing. Marketplaces often need supply and demand learning in parallel. Hardware may require early supplier, tooling and working-capital evidence.

Market sizing. Bottom-up sizing is recommended as a decision discipline, not a forecasting technique. It forces the founder to state reachable customer counts, plausible prices and channels instead of asserting an arbitrary share of a large published market.

Commitments. Deposits, pre-orders, letters of intent and pilots have different evidential strength depending on refundability, legal force, internal

effort, budget and organisational approval. The manuscript therefore asks what the customer had to give up and which threshold was crossed.

Metrics. Activation, retention, churn, customer acquisition cost and lifetime value have no useful universal benchmark across business models. The book links metrics to the natural frequency and mechanism of the product and warns against extrapolating long lifetime value from short cohorts.

Cash and runway. The book separates contribution economics from cash timing because inventory, payment terms, hiring and implementation can make growth cash-consuming even when each completed transaction has positive contribution. Detailed accounting, tax and treasury remain outside this book's scope.

Sources for the corrections and applications

Startup failure rates. The 38.4 per cent UK five-year survival figure is the current official ONS statistic for the 2019 cohort. ONS also warns about data limitations including multiple business registrations and provisional recent death estimates. The correction does not substitute a new universal startup-failure percentage.

Safety and law. The claim that safety, security, compliance and data protection can be part of minimum viability is a normative and operational principle rather than a statistic. The exact legal obligations depend on jurisdiction and industry. Founders in regulated sectors require specialist professional advice.

Finance and equity. Y Combinator materials are used only for common venture practice such as vesting and SAFEs, not as universal legal advice. Company law, tax, securities treatment and employment arrangements vary by jurisdiction.

Current verification. Current UK business-demography, financing and technology-transfer claims were checked against ONS, GOV.UK and British Business Bank materials on 9 August 2026. Current Y Combinator Startup Library materials were also checked on that date for practitioner terminology

and examples.

Bibliography

Primary and practitioner sources

Blank, Steve. "What's A Startup? First Principles." Steve Blank, 25 January 2010.

Blank, Steve, and Bob Dorf. *The Startup Owner's Manual: The Step-by-Step Guide for Building a Great Company*. Pescadero, CA: K&S Ranch, 2012.

Gascoigne, Joel. "Idea to Paying Customers in 7 Weeks: How We Did It." Buffer, 16 February 2011.

Gascoigne, Joel. "Reflecting on 10 Years of Building Buffer." Buffer, 30 November 2020.

Ries, Eric. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. New York: Crown Business, 2011.

Vohra, Rahul. "How Superhuman Built an Engine to Find Product Market Fit." *First Round Review*. Accessed 9 August 2026.

Wasserman, Noam. *The Founder's Dilemmas: Anticipating and Avoiding the Pitfalls That Can Sink a Startup*. Princeton, NJ: Princeton University Press, 2012.

Y Combinator. "How to Plan an MVP"; "The Real Product-Market Fit"; "How to Split Equity Among Co-Founders"; "Understanding SAFEs and Priced Equity Rounds." Startup Library. Accessed 9 August 2026.

Research and official sources

British Business Bank. "A Guide to Equity Funding Stages for Your Business." Business Guidance. Accessed 9 August 2026.

Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, and Chiara

Spina. "A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial." *Management Science* 66, no. 2 (2020): 564-586.

Choi, Joonkyu, Nathan Goldschlag, John Haltiwanger, and J. Daniel Kim. "Early Joiners and Startup Performance." *The Review of Economics and Statistics* 107, no. 6 (2025): 1485-1500.

Gompers, Paul A., Will Gornall, Steven N. Kaplan, and Ilya A. Strebulaev. "How Do Venture Capitalists Make Decisions?" *Journal of Financial Economics* 135, no. 1 (2020): 169-190.

Government Office for Technology Transfer. *The Knowledge Asset Spinouts Guide*. GOV.UK, updated 9 May 2025.

Kerr, William R., Ramana Nanda, and Matthew Rhodes-Kropf. "Entrepreneurship as Experimentation." *Journal of Economic Perspectives* 28, no. 3 (2014): 25-48.

Kirtley, Jacqueline, and Siobhan O'Mahony. "What Is a Pivot? Explaining When and How Entrepreneurial Firms Decide to Make Strategic Change and Pivot." *Strategic Management Journal* 44, no. 1 (2023): 197-230.

Lee, Saerom (Ronnie), and J. Daniel Kim. "When Do Startups Scale? Large-Scale Evidence from Job Postings." *Strategic Management Journal* 45, no. 9 (2024): 1633-1669.

Office for National Statistics. *Business Demography, UK: 2024*. Statistical bulletin, 20 November 2025.

Robb, Alicia M., and David T. Robinson. "The Capital Structure Decisions of New Firms." *The Review of Financial Studies* 27, no. 1 (2014): 153-179.