The scientific method in business

A way to navigate uncertainty, especially the unknown unknowns

By Dr Brian Russell

In science we start with a **hypothesis** and then ask a question to determine if the hypothesis holds up. The scientific method has a structure: **hypothesis, question, experiment, discussion, and summary of learnings**. That structure is useful in business for the same reason it's useful in science: it turns vague uncertainty into a clear process for learning what's true.

Business problems can look messier, but they're still decisions under uncertainty. They show up as: *How should we spend the budget? What's the most important goal this month, quarter, or year? Which market should we enter? Which product direction should we bet on?* The surface area is bigger, the data is noisier, and people bring strong opinions. A scientific framing helps because it gives you a shared language for turning debate into discovery, and it makes your reasoning legible so others can improve it.

The twist is that business isn't a closed system. You're operating inside a market that reacts to your choices, and you're often building something new enough that many outcomes can't be known in advance. That's where this approach becomes more than a tidy structure. It becomes an operating system for navigating uncertainty, including the kind you can't research your way into understanding.

The core idea: the only proof is a sale

There's a useful truth to hold in your mind: **until a customer exchanges dollars for your product, everything else is a probabilistic statement.** Interviews, surveys, "interest," inbound messages, pilots, waitlists, even signed letters of intent — all of it provides information, sometimes very valuable information — but it's still evidence about likelihood, not proof of value.

This isn't negative; it's clarifying. Markets are full of false positives. People are polite. People are aspirational. People want to help. People will say "yes" to an idea that sounds good and still not buy it. A sale is different because it forces a real trade: money leaves their account and enters yours. That exchange is the cleanest signal you can get that you're delivering something they value more than the alternatives.

Once you accept that, the scientific method in business becomes a way to steadily convert probability into confidence. Not perfect certainty — business never gives you that — but steadily better decisions.

The scientific method (business edition)

You can run the business version of the scientific method with a simple loop:

1. **What we know**
2. **What gap exists in what we know**
3. **What question would inform the unknown**
4. **How we will test the question**
5. **What were the results, and what did we learn**

The power isn't in the list. It's in writing it down. When it's visible, people can point to assumptions, challenge weak claims, propose better tests, and align on what "success" would look like. Strategy becomes a living document rather than a collection of opinions.

Now let's add the piece that matters most in startups and innovation: **unknown unknowns**.

What the scientific method does for unknown unknowns

Unknown unknowns aren't "we don't know the answer." They're "we don't yet know the question." They're the surprises you only discover when the product touches reality: a procurement rule you didn't anticipate, a workflow constraint, an integration requirement, a buyer persona you didn't realize exists, a compliance barrier, a distribution channel that behaves differently than expected, or a customer use-case that becomes the real product.

You can't research your way out of unknown unknowns. You can only design your process so they surface early and cheaply, before you've committed the entire company to a path.

That's why in business the scientific method isn't about being slow. It's about being deliberate: fast learning loops, clear hypotheses, and the ability to update your plan without drama.

Step 1: What we know (and what we only think we know)

Most teams accidentally blend facts, interpretations, and hopes into one narrative. The scientific habit is to separate them.

What we know in business usually looks like: actual sales, conversion rates, renewal rates, churn reasons, usage behavior, support volume, sales cycle length, average deal size, and the types of customers who repeatedly succeed with the product. What we *think* we know is everything else: that a channel will scale, that a segment will pay, that a feature will retain users, that pricing is "reasonable," that switching costs will protect you, that your product is "better."

A simple practice that changes team conversations is to label statements as one of three types: **facts, assumptions, or interpretations.** It sounds basic, but it stops a lot of strategy arguments, because many disagreements are really about which category a statement belongs in.

Step 2: Name the gap that actually decides the outcome

When teams say "we need more data," they often mean one of two things. Either they don't know what's true (do customers really have this problem, will they pay, does it deliver outcomes), or they don't know what matters most (is the bottleneck acquisition, activation, retention, or pricing; is the risk technical feasibility or distribution).

This step is where you stop collecting "more information" and start identifying the one uncertainty that matters. A good gap statement is specific and decision-relevant. It should answer: *If we resolved this, would we make a different decision?*

Step 3: Turn the gap into an executable question

Business questions are often asked in a form that invites endless debate: *"Should we enter enterprise?" "Should we raise?" "Should we build feature X?"* The scientific version turns those into questions that can be tested.

A good business question is:

- **Specific:** about a concrete outcome
- **Testable:** you can run a probe
- **Decision-relevant:** the answer changes what you do next

Instead of "Should we enter market A?" you ask, "Can we secure five qualified buyer conversations and two paid pilots in 45 days with this offer and positioning?" It forces clarity and creates a pathway to learning.

A powerful upgrade: inverse thinking

At this stage, it helps to borrow **inverse thinking**, often associated with Charlie Munger: instead of only asking "What would prove this is true?" also ask **"What would prove this is not true?"**

Disconfirming evidence is often easier to find and faster to detect. A single strong counterexample can save months of work. If your hypothesis is "Mid-market buyers will pay $20k/year for this," the inverse question is "What would show they won't?" Maybe procurement requires a bundled vendor, maybe the budget owner isn't who you thought, maybe the pain is real but not prioritized, maybe adoption requires integration that's too costly.

Finding that early isn't a failure; it's a high-quality discovery.

Step 4: Design the experiment (and remember: time is money)

In business, an experiment is not a full rollout. It's a designed probe.

The goal of an experiment is to **buy information** at the lowest possible cost in time, money, and attention. That word "time" matters more than most teams admit. **Time is money** in startups. If you experiment too long, you pay with calendar time, momentum, team energy, and competitive positioning. Over-experimentation can be more expensive than building, because it delays the moment where reality starts talking back.

So the trick is to balance two truths:

- Research reduces predictable mistakes and avoids dead ends that violate physics, regulations, or basic economics.
- Execution reveals hidden constraints and hidden opportunities, especially unknown unknowns, that research can't uncover.

This means "research vs execution" is not a debate. It's a curve. Every decision sits somewhere on it. High downside and irreversible outcomes push you toward more research first. Low downside and fast feedback push you toward execution earlier.

A practical way to choose your point on the curve is to ask:

- What is the cost of being wrong?
- Is it reversible or irreversible?
- How quickly will execution produce a clear signal?
- Can we design a test that is cheap, fast, and safe?

When the cost of being wrong is low and the feedback is fast, moving quickly is rational. When the cost of being wrong is high or irreversible, you buy more certainty first.

The hard case: when you need millions and years before you can sell anything

Some of the most meaningful startups don't get the luxury of fast tests. Hardware, regulated industries, deep tech, infrastructure, and complex integrations can require **millions of dollars and years** before the product can truly be sold. In these cases, "research" often becomes indistinguishable from "building." You can read more, interview more, model more — but certain truths only appear when something exists in the world.

In long-cycle startups you are running **two hypotheses in parallel** from day one.

The first is the **problem hypothesis:** customers agree the problem exists, it matters, and it's worth solving now.

The second is the **solution hypothesis:** your product will be the way customers actually want to solve that problem inside their real constraints.

These can be independently true. Customers can passionately agree the problem exists and still reject your solution because adoption is hard, integration is expensive, procurement is slow, compliance requirements are heavy, or the buyer's incentives are misaligned.

So in long-cycle startups, the scientific method becomes: **how do we reduce these hypotheses into smaller risks without waiting years for the final sale?**

Risk vs perceived risk

This is where it helps to separate **risk** from **perceived risk**.

- **Risk** is what is actually true about feasibility, adoption, economics, and competition.
- **Perceived risk** is what customers and investors believe is risky, based on what they can verify today and what their incentives reward.

Large companies and professional investors often want more proof not because they dislike innovation, but because their job is to avoid avoidable downside. Proof compresses perceived risk. The earlier you are, the less final proof you can provide, so you build a ladder of evidence that steadily reduces perceived risk enough to unlock the next step.

In a long-cycle startup, early milestones are often about proving credibility and reducing uncertainty: prototypes, benchmark tests, simulations, pilot structures, design partnerships, regulatory milestones, manufacturing feasibility, integration proofs, and early outcome demonstrations. None of these are the final proof that a sale provides, but each reduces risk or perceived risk and makes the next commitment rational.

Intuition can be right, even without large evidence

There's another reality in entrepreneurship: sometimes founders are right because they have **earned intuition.** They've lived inside the domain. They understand the workflow, the incentives, the supply chain, the buying process, and the failure modes. They can see the shape of the future earlier than the market can articulate it.

This is why entrepreneurs can be right without massive evidence or perfect theoretical research. Many markets can't clearly describe what they will value until a solution exists. Early customer feedback is often biased toward incremental improvements because that's what people can confidently imagine.

The scientific method doesn't oppose intuition. It refines it. The best pattern is: intuition proposes the hypothesis, then you design experiments that either support it or quickly reveal what needs to change.

Step 5: Discussion and learnings (results only matter if they change the plan)

After the experiment, don't just report what happened. Update what you know. That closes the loop and creates compounding advantage.

Over time you build a library of proven truths about your market: what people pay for, why they churn, which message converts, which segment expands, which channel scales, which integrations are mandatory, which features are distractions. That knowledge becomes a durable asset because it's expensive for competitors to recreate quickly.

This is also where humility becomes a strategy. Instead of "this will work," the language becomes "based on evidence, this has a high likelihood, and here's the next test to increase confidence." That mindset keeps teams adaptive without being indecisive.

A practical tie-breaker when you can't afford more research

Sometimes you reach a point where the next increment of "research" is really just building the product. You won't learn the remaining truths by debating. You'll learn them by shipping something real.

That's also when teams hit a common situation: two smart people propose two different approaches. Both are plausible. Neither can be proven better with the data you have. The argument produces heat, not information.

A simple tie-breaker that protects speed is this:

Choose the approach most dearly held by one team member.

Not because it's automatically correct, but because strong belief often produces faster, more committed execution. When someone feels genuine ownership, they'll drive it through rough edges and take responsibility for making it work. In early-stage work, where both approaches are "reasonable," the biggest variable is often not the approach — it's the quality and speed of implementation.

Keep it scientific by adding a final clause:

Commit hard, instrument the build, and be willing to revisit if evidence disagrees.

That creates alignment and speed without becoming stubborn.

A worked example: “How should we spend the budget?”

Budget debates are often disguised bottleneck debates. Scientific framing turns them into a learning cycle.

Start with what you know. Perhaps paid ads drive traffic but conversion is weak, or outbound converts well once prospects understand the category, or retention is strong after activation. Name the gap: is the constraint top-of-funnel volume, conversion, or retention? Then define a testable question: “If we improve onboarding and messaging clarity, does conversion rise enough that the same spend produces more revenue?”

Run a two-week probe: change onboarding and messaging while holding spend constant. Measure conversion from visit to signup to first value to paid. If conversion improves, you’ve earned the right to spend more on acquisition. If it doesn’t, you’ve learned that more spend won’t fix the system yet, and the next experiment should target product, offer, or segment.

This is how the scientific method converts “opinions” into “evidence,” quickly.

The simplest weekly cadence

If you want this to become a habit, run it as a weekly thread:

What we know now.
What we don’t know yet.
The question that matters most.
The smallest test that reveals the truth (including what would disprove it).
What changed because of the result.

That’s the scientific method in business. It respects the reality that the only proof is a sale, it makes space for unknown unknowns, it balances learning with the fact that time is money, and it gives intuition a structure that turns vision into progress.