INQUIRY · INQUIRY IN ORGANISATIONS

Hypothesis-Driven Business Strategy

Every business strategy rests on beliefs about customers, competitors and capabilities that may be wrong. This article explains how to treat strategy as a set of hypotheses to be tested, what research on discovery-driven planning and scientific decision making shows, and the practical steps involved.

ModQu research article · 2,929 words · 26 references (APA 7th edition)

Every strategy is a bet. A company that enters a new market is betting that customers there want what it offers, that it can reach them at a reasonable cost and that competitors will not respond in a way that eliminates the profit. A company that cuts prices is betting that the extra volume will outweigh the lower margin. These bets rest on beliefs about the world, and some of those beliefs will be wrong. Hypothesis-driven strategy makes the bets explicit. It treats a strategy as a set of hypotheses about what will create value, identifies the assumptions on which they depend and tests the most important and uncertain ones before committing large resources. This article explains the approach, its roots in research on planning and decision making, and the evidence that it improves results.

Strategy as theory

Several strategists have argued that a strategy is best understood as a theory. Richard Rumelt (2011) described a good strategy as built on a diagnosis of the situation, a guiding policy for dealing with it and coherent actions to carry out the policy. He emphasised that the diagnosis and policy amount to a hypothesis about what will work, which must be judged and revised in the light of results. Robert Kaplan and David Norton (1996), in developing the balanced scorecard, described strategy as a set of hypotheses about cause and effect: for example, that training employees will improve service quality, which will increase customer loyalty, which will raise revenue. Their scorecard was designed to make these links explicit so that they could be measured and tested.

Teppo Felin and Todd Zenger (2017) proposed a *theory-based view* of strategy. They argued that successful strategies often depend on a distinctive theory held by managers about how value can be created, a theory that others do not share and that cannot be derived from existing data alone. Such theories are valuable precisely because they are not obvious. But because they are not obvious, they must be tested, and the testing should be designed to probe the theory's specific claims. Zenger (2013) suggested that every firm should be able to state its own theory: what it believes about its industry, its customers and its capabilities that justifies the choices it makes.

Lafley et al. (2012), drawing on strategy work at Procter & Gamble, proposed bringing scientific method into strategic choice. Instead of debating which option is best, teams should ask what would have to be true for each option to succeed, and then design tests of the conditions they are least sure about. This turns arguments about preferences into questions that evidence can settle.

Deliberate and emergent strategy

Henry Mintzberg and James Waters (1985) distinguished *deliberate* strategies, which are planned in advance and realised as intended, from *emergent* strategies, which arise from patterns of action that were not planned. In practice, most realised strategies are a mixture. Some planned elements fail and are abandoned; some unplanned initiatives succeed and are adopted. Hypothesis-driven strategy can be seen as a way of managing this mixture deliberately: planning what can be planned, while setting up the tests that will reveal where the plan is wrong and what should emerge in its place.

Discovery-driven planning

Rita McGrath and Ian MacMillan (1995) argued that conventional planning is poorly suited to new ventures. Conventional plans assume that the future can be extrapolated from a base of knowledge, and they measure success by whether results meet the plan. For a new venture, the ratio of assumptions to knowledge is high, and a plan that meets its targets on paper may rest on assumptions that have never been checked.

McGrath and MacMillan proposed *discovery-driven planning*. It begins with the required outcome, such as the profit the venture must earn to be worthwhile, and works backwards to determine what would have to be true for that outcome to be achieved: how many customers, at what price, with what costs. Each of these becomes an explicit assumption. The plan then sets out *checkpoints*, milestones at which specific assumptions will be tested before further resources are committed. If the tests show that key assumptions are false, the plan is revised or abandoned while the cost of doing so is still small. McGrath and MacMillan (2009) later extended the approach to corporate growth initiatives more generally.

The research on testing assumptions

Several lines of research support the value of explicitly stating and testing strategic assumptions.

Structured debate. Mason and Mitroff (1981) developed *strategic assumption surfacing and testing*, in which groups identify the stakeholders and assumptions underlying competing strategies, rate each assumption for importance and certainty, and debate the assumptions rather than the conclusions. Figure 1 shows how this rating directs attention.

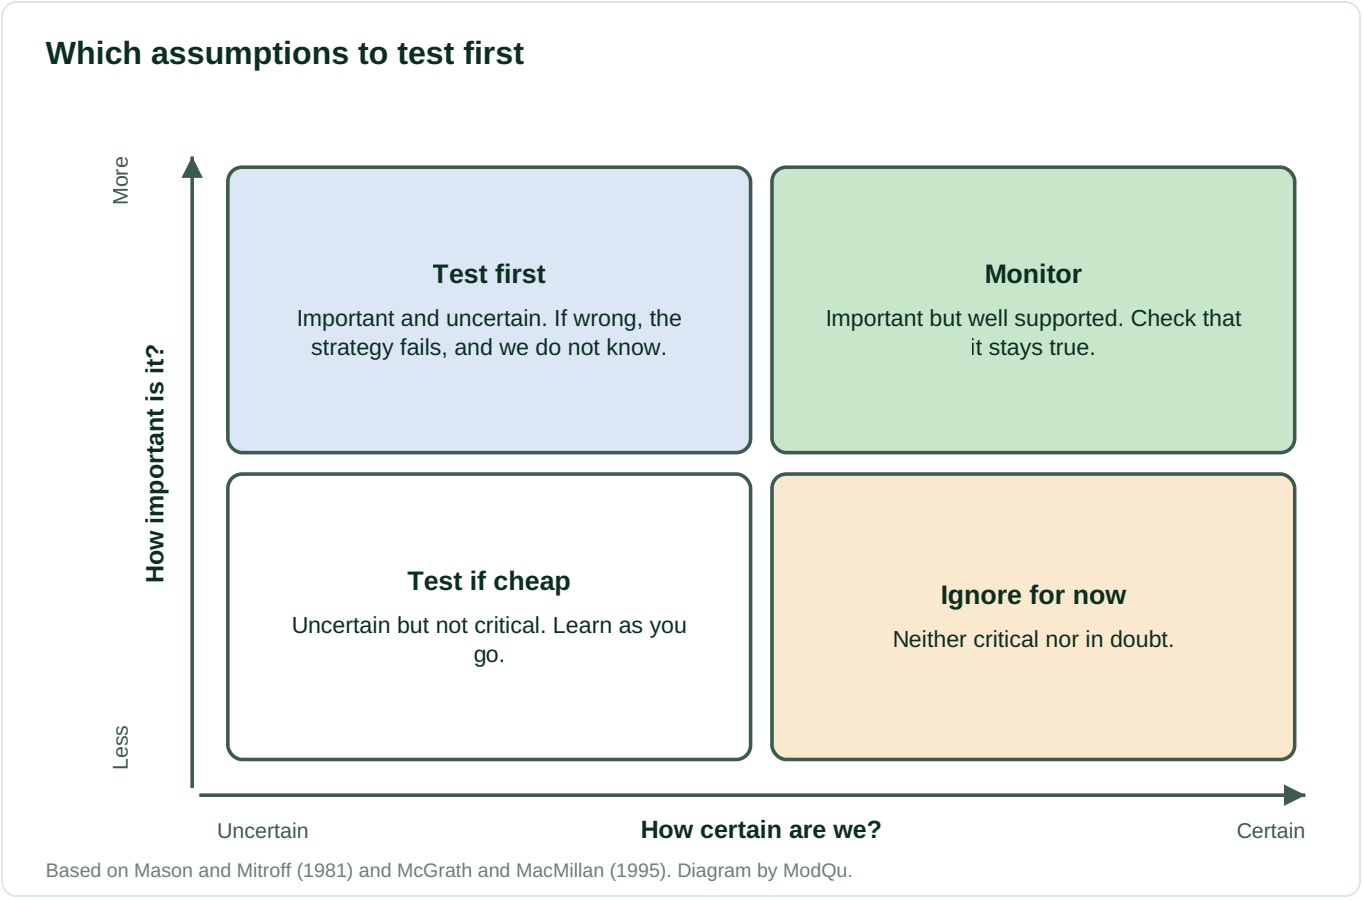


Figure 1. Prioritising assumptions by importance and certainty. The assumptions to test first are those that matter most and are least certain. Based on Mason and Mitroff (1981) and McGrath and MacMillan (1995); diagram by ModQu.

Overconfidence in plans. Daniel Kahneman and Dan Lovallo (1993) showed that managers tend to make "bold forecasts", taking an *inside view* of their project that focuses on its specific plans and scenarios while neglecting the *outside view*: the record of similar projects. Lovallo and Kahneman (2003) argued that this produces systematic optimism in business plans, with costs underestimated and benefits overestimated. Treating plans as hypotheses, and comparing them with the outcomes of similar ventures, is a direct remedy.

Confirmation bias. Raymond Nickerson (1998) reviewed extensive evidence that people tend to seek, interpret and remember information in ways that confirm their existing beliefs. Klayman and Ha (1987) showed that people typically test hypotheses by looking for cases where they expect the hypothesis to hold, a strategy that can fail to reveal when the hypothesis is wrong. In strategy, this means that managers often look for evidence that their plan will work rather than for evidence that it will not. A hypothesis-driven approach asks explicitly what result would show the strategy to be wrong, in the spirit of Karl Popper's (1959) argument that a claim is tested only when it is exposed to possible refutation.

Escalation of commitment. Barry Staw (1976) found that people who were personally responsible for a failing investment tended to commit more resources to it, rather than less, in an effort to justify their earlier decision. Setting explicit checkpoints and criteria for stopping in advance makes it easier to abandon a strategy whose assumptions have failed.

Evidence that a scientific approach helps

The most direct evidence comes from a randomised trial. Camuffo et al. (2020) worked with 116 Italian start-ups. Half were trained to make decisions in a scientific way: to formulate their business ideas as theories, derive explicit hypotheses, test them with rigorous experiments and evaluate the results objectively. The other half received training of similar length on the same topics without this emphasis. Firms trained in the scientific approach were more likely to pivot to a different idea or to abandon their project, and those that continued earned more revenue. The researchers argued that the scientific approach helped entrepreneurs recognise more accurately when their ideas were not working and search more effectively for alternatives.

William Kerr and colleagues (2014) described entrepreneurship as a process of experimentation, in which the main uncertainty is whether an idea will work and the value of early experiments lies in revealing that information cheaply. Gans et al. (2019) argued that because entrepreneurs face choices among many possible strategies that cannot all be tested, the order and design of experiments is itself a strategic decision.

In established firms, Stefan Thomke (2020) documented how companies that run large numbers of controlled experiments, testing changes to products, prices and processes, make better decisions than those that rely on judgement alone. Kohavi et al. (2020) report that in large online businesses, a majority of ideas fail to improve the measures they were intended to improve when tested in controlled experiments. The finding is a reminder that strategic intuition, however experienced, is often wrong about specifics.

Matching tests to assumptions

Different kinds of strategic assumption call for different kinds of evidence. Choosing the right test is as important as deciding to test at all.

Assumptions about customer needs can often be explored through interviews and observation, which reveal what problems customers face and how they currently solve them. These methods are good at discovering needs that were not anticipated, but they involve small numbers and cannot show how common a need is.

Assumptions about market size and demand call for evidence from larger, representative samples. Surveys can estimate how many potential customers share a need and how strongly they feel it. Questions about willingness to pay need careful design, because stated intentions tend to overstate actual behaviour. Asking people to choose between realistic alternatives, rather than to state a price, usually gives more reliable estimates.

Assumptions about behaviour are best tested by observing what people actually do. A pilot launch in one region, a landing page that measures how many visitors sign up for a product not yet built, or a controlled experiment comparing two offers all provide behavioural evidence. Thomke (2020) argues that controlled experiments are the most reliable way of establishing cause and effect in business, because they separate the effect of a change from everything else happening at the same time.

Assumptions about competitors are harder to test directly, because competitors' responses only occur once the strategy is visible. Evidence can come from studying how competitors responded to similar moves in the past, from analysing their capabilities and incentives, or from structured exercises in which a team plays the role of the competitor and tries to defeat the strategy. The outside view recommended by Kahneman and Lovallo (1993), which asks how similar initiatives have fared, is especially useful here.

Assumptions about the organisation's own capabilities can be tested through small internal trials. Can the sales team actually reach the new customer group? Can operations deliver at the required cost? These assumptions are often treated as certain because they concern the company itself, yet they are frequently where strategies fail.

In each case, the test should be designed so that it could produce a disappointing answer. A survey that only asks satisfied customers, or a pilot run in the most favourable market, may confirm the strategy without testing it.

The lean start-up

The best-known popular form of hypothesis-driven strategy is the *lean start-up*. Eric Ries (2011) proposed that start-ups should treat their business model as a set of hypotheses, build a *minimum viable product* to test the most important of them and use the results to decide whether to persevere or pivot. Steve Blank (2013) argued that the approach reverses the traditional sequence of writing a business plan and then executing it. Instead, founders go out and test their hypotheses with customers, and the business model is revised as evidence accumulates. The approach has been widely adopted by large companies for internal ventures as well as by start-ups.

Not all strategy can be tested in advance

Hypothesis-driven strategy has limits. Some strategic choices cannot be tested on a small scale because their effects depend on scale itself, or on competitors' responses that will only occur once the strategy is visible. Some opportunities are fleeting, and the time taken to test may allow competitors to act first. Some hypotheses can only be tested by committing substantial resources.

Saras Sarasvathy (2001) observed that expert entrepreneurs often do not begin with hypotheses about a target market at all. Instead, they start from the means they have and allow goals to emerge through action and partnerships. Her *effectuation* approach suggests that in highly uncertain settings, controlling what can be controlled may be more useful than predicting what cannot. Hypothesis-driven strategy and effectuation can be complementary: effectuation for exploring possibilities when almost nothing is known, hypothesis testing for deciding which possibilities to pursue once candidate strategies exist.

There is also a risk of testing what is easy rather than what matters. A team may run many experiments on details such as website wording while leaving the central assumptions of the strategy untested. Prioritising by importance and uncertainty is the guard against this.

A practical method

The research suggests a practical sequence, shown in Figure 2.

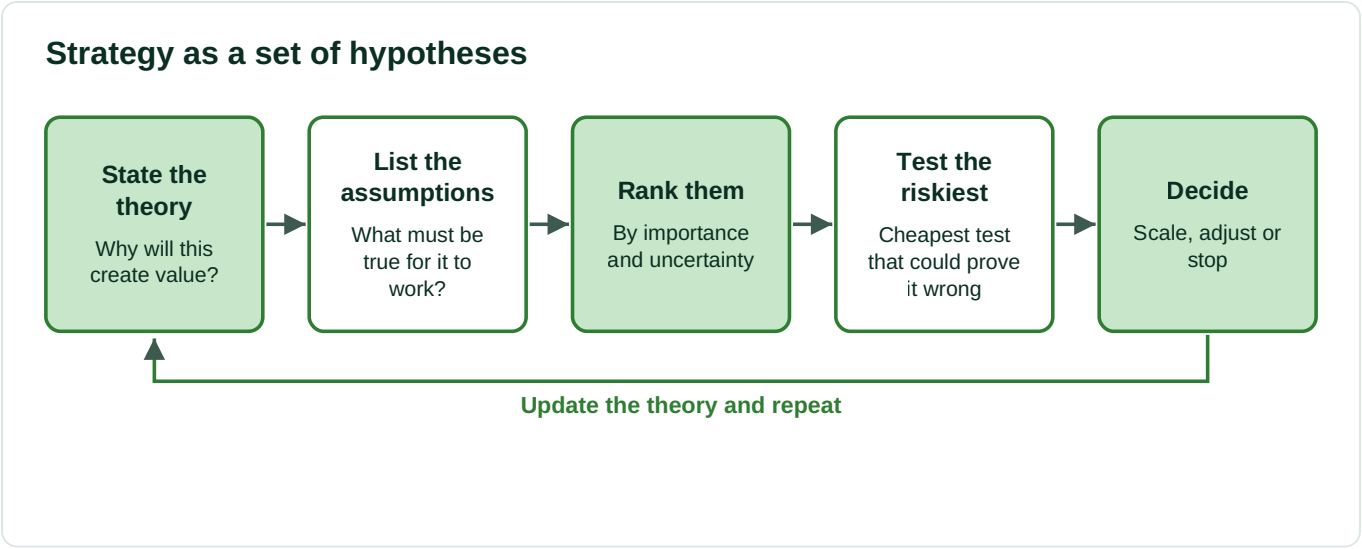


Figure 2. Strategy as a set of hypotheses to be tested. Diagram by ModQu.

1. State the theory. Write down, in a few sentences, why the strategy will create value. Who will buy, why, and why will the company be able to serve them better than alternatives?

2. List the assumptions. For each element of the theory, ask what must be true. Work backwards from the required outcome, as discovery-driven planning recommends: how many customers, at what price, with what costs, on what timetable?

3. Rank the assumptions. Rate each for importance (would the strategy fail if it were wrong?) and certainty (how good is the current evidence?). Focus on those that are both important and uncertain.

4. Design the cheapest test that could prove each assumption wrong. A test might be a survey of target customers, a small pilot, a controlled experiment, an analysis of comparable cases or a set of interviews. Decide in advance what result would count as disconfirming. Gary Klein's (2007) *pre-mortem*, in which the team imagines that the strategy has failed and explains why, is a useful way to find assumptions that have been overlooked.

5. Decide. On the basis of the results, scale the strategy, adjust it or stop. Record what was learned and update the theory.

6. Repeat. Strategy execution generates new evidence continuously. Sull et al. (2015) found that a common reason for poor strategy execution is a failure to adapt to changing conditions, often because organisations stick rigidly to plans. Building regular reviews of assumptions into the strategic process allows adaptation to be deliberate.

Expressing hypotheses well

A good strategic hypothesis is specific enough to be tested. "Customers want better service" cannot easily be disproved. "At least 30% of our small business customers would pay a 10% premium for same-day support" can be. A useful format states the expected result, the population, the measure and the threshold that would count as success.

Hypotheses should also be stated with an honest level of confidence. Tetlock and Gardner (2015) showed that forecasters who expressed their beliefs as specific probabilities, and updated them as evidence arrived, were more accurate than those who used vague language. A strategy team that says "we think there is a 60% chance that this market will grow by more than 5% a year" can learn from the outcome in a way that a team saying "the market looks promising" cannot.

A hypothetical example

A software company that sells accounting tools to small businesses is considering a new product for mid-sized firms. The strategy team writes down its theory: mid-sized firms are poorly served by existing products, which are either too simple or too complex, and the company can adapt its existing tool to fill the gap at modest cost.

Working backwards from the revenue the new product must earn to justify the investment, the team lists the assumptions: that at least 2,000 mid-sized firms in its markets are dissatisfied with their current tools; that they would switch for a product meeting their needs; that they would pay at least twice the small business price; that the existing tool can be adapted with six months of development; and that the company's sales channels, which currently serve small businesses, can reach mid-sized buyers.

The team rates these assumptions. The size of the dissatisfied market and the willingness to switch are important and uncertain. The development time is important but relatively well understood. The sales channel assumption is important and highly uncertain, because the company has never sold to larger firms. The team designs three tests: a survey of mid-sized firms to estimate dissatisfaction and switching intentions; a set of interviews with finance managers to understand how they buy software; and a small trial in which the existing sales team tries to arrange meetings with mid-sized prospects.

The survey supports the market assumption. The interviews reveal that mid-sized firms buy through a formal process involving several decision makers, which the existing sales team is not set up to handle. The trial confirms this: few meetings are arranged. The team revises the strategy to include a partnership with a reseller that already serves mid-sized firms, and tests that assumption next. The original plan would have failed on the sales channel assumption, at much greater cost, had it been executed without testing.

Conclusion

Hypothesis-driven strategy applies the logic of scientific inquiry to business decisions. It does not remove the need for judgement or vision; the theory at the heart of a strategy may be bold and original. What it adds is discipline in identifying what that theory assumes, finding out which assumptions are weakest and testing them before the cost of being wrong becomes large. Research on overconfidence, confirmation bias and escalation of commitment explains why this discipline does not come naturally. Research on discovery-driven planning and scientific decision making shows that it can be learned and that it improves results. For organisations that use research to inform strategy, the practical implication is clear: the most valuable research is often the study that tests the assumption the strategy depends on most.

References

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

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science, 66*(2), 564–586. https://doi.org/10.1287/mnsc.2018.3249

Felin, T., & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science, 2*(4), 258–271. https://doi.org/10.1287/stsc.2017.0048

Gans, J. S., Stern, S., & Wu, J. (2019). Foundations of entrepreneurial strategy. *Strategic Management Journal, 40*(5), 736–756. https://doi.org/10.1002/smj.3010

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science, 39*(1), 17–31. https://doi.org/10.1287/mnsc.39.1.17

Kaplan, R. S., & Norton, D. P. (1996). *The balanced scorecard: Translating strategy into action*. Harvard Business School Press.

Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives, 28*(3), 25–48. https://doi.org/10.1257/jep.28.3.25

Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*(2), 211–228. https://doi.org/10.1037/0033-295X.94.2.211

Klein, G. (2007). Performing a project premortem. *Harvard Business Review, 85*(9), 18–19.

Kohavi, R., Tang, D., & Xu, Y. (2020). *Trustworthy online controlled experiments: A practical guide to A/B testing*. Cambridge University Press. https://doi.org/10.1017/9781108653985

Lafley, A. G., Martin, R. L., Rivkin, J. W., & Siggelkow, N. (2012). Bringing science to the art of strategy. *Harvard Business Review, 90*(9), 56–66.

Lovallo, D., & Kahneman, D. (2003). Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review, 81*(7), 56–63.

Mason, R. O., & Mitroff, I. I. (1981). *Challenging strategic planning assumptions: Theory, cases, and techniques*. Wiley.

McGrath, R. G., & MacMillan, I. C. (1995). Discovery-driven planning. *Harvard Business Review, 73*(4), 44–54.

McGrath, R. G., & MacMillan, I. C. (2009). *Discovery-driven growth: A breakthrough process to reduce risk and seize opportunity*. Harvard Business Press.

Mintzberg, H., & Waters, J. A. (1985). Of strategies, deliberate and emergent. *Strategic Management Journal, 6*(3), 257–272. https://doi.org/10.1002/smj.4250060306

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.

Ries, E. (2011). *The lean startup*. Crown Business.

Rumelt, R. P. (2011). *Good strategy/bad strategy: The difference and why it matters*. Crown Business.

Sarasvathy, S. D. (2001). Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review, 26*(2), 243–263. https://doi.org/10.5465/amr.2001.4378020

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance, 16*(1), 27–44. https://doi.org/10.1016/0030-5073(76)90005-2

Sull, D., Homkes, R., & Sull, C. (2015). Why strategy execution unravels—and what to do about it. *Harvard Business Review, 93*(3), 57–66.

Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Crown.

Thomke, S. H. (2020). *Experimentation works: The surprising power of business experiments*. Harvard Business Review Press.

Zenger, T. (2013). What is the theory of your firm? *Harvard Business Review, 91*(6), 72–78.