


Innovation Management

From scattered initiatives
to a sustained organizational capability

Dr. Omar A. Alsamani



Innovation Management

From scattered initiatives
to a sustained organizational capability

Dr. Omar A. Alsamani

Innovation Management

From scattered initiatives to a sustained organizational capability

This book is published under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International licence (CC BY-NC-ND 4.0). It may be copied, distributed and shared, in full or in part, in any medium or format, for non-commercial educational and research purposes, provided it is attributed to its author and its text is not altered. Licence text: creativecommons.org/licenses/by-nc-nd/4.0/

Website: alsamani.com

© 2026 Omar A. Alsamani

First edition, 1448 AH / 2026

Legal deposit no. 1448/7324

ISBN 978-603-06-7961-4

14 Experimentation and risk reduction

Between an idea and a major investment lies a stretch of assumptions, and experimentation is what turns them into knowledge before the organization commits to something that is hard to reverse. Within this horizon we move step by step from ranking assumptions by how critical they are, to the types of tests available, to designing an experiment that yields genuine learning, and we close with the common mistakes made in pilots.

Introduction

Opening case: An appointment system (composite case)

A service organization developed a new appointment booking system for its branches and chose its main branch in the capital to pilot it: the newest branch, the most heavily staffed, and the closest to the project team. Three extra staff members were assigned to the pilot to help visitors, and the development team came in daily to fix any problem on the spot. After three months, the pilot was declared a success: waiting times had fallen and visitor satisfaction had risen. The system was then rolled out to forty branches. In most of

them, work was disrupted for weeks: the small branches had no one to help visitors, some had slower networks, and the development team could not be everywhere at once.

The pilot did not fail; it was designed to succeed. The system was tested under the best possible conditions, so the pilot proved that it worked under those conditions and said nothing about the conditions in which it would actually operate. This difference between an experiment that proves and one that teaches is what this chapter is about.

Learning Objectives

After studying this chapter, the reader should be able to:

- 1 Explain experimentation as a means of turning assumptions into knowledge and reducing uncertainty.
- 2 Rank an initiative's assumptions by their importance and the strength of the evidence behind them.
- 3 Choose the appropriate type of test for each assumption and stage.
- ε Design an experiment with a hypothesis, a measure, and a success criterion set in advance.
- 0 Avoid the common mistakes made in pilots.

14.1 Managing uncertainty through learning

In conventional projects, the plan is built on knowledge, and success is measured by adherence to it. In innovation, by contrast, the initial plan rests more on assumptions than on knowledge. McGrath and MacMillan observed that many new ventures fail because their founders treat assumptions as if they were facts, only to discover their errors after everything has been

invested. They proposed what they called **discovery-driven planning**: writing the assumptions down explicitly and sequencing the stages so that each stage tests the most important assumptions before moving on, so that assumptions are gradually converted into knowledge (McGrath & MacMillan, 1995).

Experimental evidence supports this approach. Camuffo and colleagues conducted a randomized trial with 116 start-ups in Italy, half of which received training in making decisions by a scientific method: formulating hypotheses, testing them, and judging the results. The trained start-ups were more likely to change direction or stop when the evidence did not support their idea, and they achieved higher revenues (Camuffo et al., 2020). The experimental approach leaves the odds of the first idea succeeding unchanged; its benefit lies in knowing when the idea needs to change.

This is the logic underlying the gates of Chapter 10 and staged funding: each stage buys knowledge that reduces uncertainty, and the next stage is funded according to what the team has learned. Experimentation is the tool with which the team buys this knowledge. It is consistent with the principle of “managing uncertainty” among the principles of innovation management: dealing with risks by confronting them and learning from them.

14.2 Critical assumptions and how to rank them

Chapter 12 ended with a testable concept that states its critical assumptions, and Chapter 13 noted that the business model itself rests on assumptions. Experimentation begins by gathering all of these assumptions and then ranking them. They can be classified by the four evaluation lenses: assumptions about *desirability* (Do people want the solution?), *feasibility* (Can it be built and operated?), *viability* (Can it be sustained?), and *responsibility* (Is it fair and safe?).

Testing all the assumptions at once is neither possible nor useful. Priority goes to assumptions that combine two qualities: they are **important**, meaning that if they are wrong the entire initiative collapses, and the **evidence for them is**

weak. Bland and Osterwalder propose ranking them along these two dimensions (Bland & Osterwalder, 2019), as Figure 14.1 shows.




Figure 14.1 Ranking assumptions by their importance and the strength of the evidence behind them, adapted from Bland & Osterwalder (2019). The top row indicates high importance, and the right-hand column strong evidence.

Teams tend to start with what is easy to test before what must be tested. A technical team tests feasibility because that is what it does well, and postpones the desirability assumption because testing it means going out to people. The result is a solution that works technically and that nobody uses. Ranking assumptions in this way puts risk ahead of convenience.

14.3 The testing ladder

Tests vary in their cost and in the strength of the evidence they provide. The principle of choice is simple: start with the cheapest test that could disprove the assumption, and move up to a costlier test only if the assumption survives the cheaper one. Figure 14.2 shows a ladder of such tests.




Figure 14.2 The testing ladder, from cheapest to costliest. Prepared by the author.

Every rung on the ladder requires a **prototype** of some kind. Houde and Hill pointed out that the most important question in designing a prototype is not “What are we building?” but “What do we want to find out?” They distinguished among prototypes that test the *role* of the solution in the user's life, prototypes that test its *look and feel*, and prototypes that test its technical *implementation* (Houde & Hill, 1997). Each question has its own prototype: one that answers the feasibility question falls short on desirability, and perfecting a prototype that answers the wrong question is waste.

One of the most widespread concepts in this area is the **minimum viable product**, the smallest version of a product that allows the team to learn as much as possible about users with the least effort (Ries, 2011). The key word in the definition is “learn”: the initial version is a tool for testing an assumption, not a stripped-down version of the final product that is launched and then added to.

In Practice

Before your team asks for a budget to build an initial version, ask: Which assumption will this version test? And is there a cheaper test lower down the ladder that could disprove it? Many questions can be answered through conversations or manual delivery of the service before a single line of code is written.

14.4

Designing a learning experiment

The difference between an experiment that teaches and one that merely proves lies in the design, before the start. A good experiment needs five elements defined in advance:

- **A clear hypothesis:** “We believe that [a specific group] will do [a specific behavior] because [a reason].”
- **A behavioral measure:** one that captures what people actually do; actual use is more telling than stated opinion.
- **A success criterion set in advance:** the figure that, if reached, means the hypothesis is supported, and if not reached, means it is rejected. Setting it after seeing the results opens the door to interpreting any result as a success.
- **A comparison:** with the status quo or with an alternative, because improvement may occur for reasons that have nothing to do with the solution.

- **Duration, limits, and stopping conditions:** when the experiment ends, to whom it applies, and what would halt it early if harm emerges.

Comparison is the hardest of these elements and the most often neglected in the public sector. The UK government's Behavioural Insights Team has advocated the use of randomized controlled trials in public policy, showing that in many cases they are less costly and complex than is commonly assumed, and that they allow organizations to learn what actually works before scaling it up (Haynes et al., 2012). Where random assignment is not feasible, comparison with similar branches or periods is still far better than no comparison at all.

At a university, these elements can become part of applied projects in courses: instead of presenting a solution and defending it, students present a hypothesis and test it with real users, and they are assessed on the quality of the experiment and what they learned from it, even if their hypothesis is rejected. This links the course to the innovation system and teaches students the method before it teaches them the result.

14.5 Pilots and their pitfalls

The pilot is the highest and costliest rung on the ladder, and it is run once the core assumptions have survived the cheaper tests. Its function is to reveal how the solution works under real operating conditions; it is a tool for learning more than for proof. Certain mistakes recur in pilots:

- **Ideal conditions:** choosing the best site, the most enthusiastic staff, and exceptional support from the project team, as in the opening case. The alternative is to run the pilot at a site that represents ordinary conditions, or at two different sites.
- **The never-ending pilot:** a pilot extended again and again without a decision, which either becomes a permanent service with no owner and no budget or dies quietly. This is one form of decision limbo.

- **No future owner:** a pilot run by the innovation team alone, which finds no one to adopt it once it ends, as in the opening case of Chapter 13.
- **Measuring what is easy to measure:** the number of users or their satisfaction in a survey, rather than the behavior that signals value: Did they come back? Did completion time change? Did errors decrease?
- **Overlooking the true cost of operation:** exceptional support during the pilot hides what routine operation will require.

Every pilot has three possible outcomes, all of them useful if the pilot is well designed: the solution is scaled up, it is modified and tested again, or it is stopped. Chapter 15 addresses the move from pilot to scale.

Caution

An experiment that cannot fail teaches nothing. If the team cannot specify before the start the result that would lead it to stop or modify the initiative, then what it is running is a demonstration, even if it is called an experiment.

14.6 A culture of experimentation

Organizations that experiment at scale uncover a seemingly discouraging truth: most ideas fail. Kohavi and Thomke reported that only 10 to 20 percent of experiments at major search engines produce positive results, and that small changes whose effect no one anticipated sometimes yielded large returns, such as a tweak to the way ad headlines were displayed that raised one search engine's revenue by about 12 percent (Kohavi & Thomke, 2017). The implication is that experience and intuition, however valuable, still need testing, and that an organization implementing its ideas without testing them implements many unhelpful ideas along with them.

A culture of experimentation is broader than its tools. It requires what Chapter 8 discussed: psychological safety that allows a team to announce that its hypothesis was rejected, a distinction between intelligent failure and avoidable failure, and leadership that asks “What did we learn?” before it asks “Why did we fail?” It also requires respect for the people on whom experiments are conducted: the beneficiary of a public service is not a test subject, and has a right to be informed and protected from harm. Chapter 19 addresses the ethics of experimenting on people.

Chapter tool: The experiment card

Write a card for each experiment before it begins, and ask someone outside the team to review it. Once the experiment is over, complete the last two fields without altering anything written above them.

Chapter tool The experiment card (example from the opening case, as it should have been designed)

Item	Question	Example
Assumption tested	Which critical assumption does the experiment test?	That branches with ordinary resources can run the system without additional support
Hypothesis	We believe that... will... because...	We believe that visitors to a mid-sized branch will book their appointments without assistance because the interface is simple
Type of test	Which rung on the ladder? Why?	A pilot in two branches: one mid-sized and one small
Measure	What behavior will we measure?	The share of bookings completed without staff intervention, and waiting time

Item	Question	Example
Success criterion	What figure supports the hypothesis? (Set before the start)	Seventy percent of bookings completed without intervention, and waiting time cut by a quarter
Comparison	What do we compare against?	Two similar branches running the old system over the same period
Duration and limits	When does it end? To whom does it apply?	Six weeks, in the two branches only, with no additional staff
Stopping conditions	What would halt the experiment early?	Waiting time doubling for two consecutive weeks
Future owner	Who will adopt the solution if it succeeds? Are they involved?	The branch management department, represented on the pilot team
Result	What happened? (After the experiment)	
What we learned and the decision	Scale up, modify, or stop?	

Chapter Summary

- In innovation, the plan rests on assumptions, and experimentation turns them into knowledge before major investment; the most valuable thing about the experimental approach is that it reveals when an idea needs to change.
- Assumptions are ranked by their importance and the strength of the evidence behind them, with priority given to important assumptions with weak evidence.

- Start with the cheapest test that could disprove the assumption, and move up the ladder if it holds. A prototype is designed to answer a specific question.
- A learning experiment needs a hypothesis, a behavioral measure, a success criterion set in advance, a comparison, and limits and stopping conditions.
- Common mistakes in pilots include ideal conditions, the never-ending pilot, the absence of an owner, and measuring what is easy to measure.
- Most ideas fail when tested, which is why testing remains necessary however strong the intuition. A culture of experimentation requires psychological safety and respect for those on whom experiments are conducted.

Key Terms

Discovery-driven planning Discovery-driven planning	Planning that writes assumptions down explicitly and sequences stages to test the most important ones first.
Critical assumption Critical assumption	An assumption that, if wrong, brings down the entire initiative.
Prototype Prototype	A representation of the solution built to answer a specific question about its role, its look and feel, or its implementation.
Minimum viable product Minimum viable product	The smallest version of a product that enables the most learning with the least effort.

Randomized controlled trial Randomized controlled trial	An experiment in which participants are randomly assigned to the solution or an alternative, in order to isolate the solution's effect.
Pilot Pilot	Implementing the solution under real operating conditions, within clear limits, before scaling it up.

Review and Discussion Questions

- 1

What is the difference between an experiment that proves and one that teaches? Illustrate your answer with the opening case.

Understanding
- 2

Write down five assumptions for an initiative you are working on, and rank them on Figure 14.1.

Application
- 1

Take the top-ranked assumption and identify the cheapest test that could disprove it.

Application
- 2

Analyze the most recent pilot in your organization in light of the five mistakes in Section 14.5.

Analysis
- 3

“Randomized controlled trials are a luxury ill-suited to government agencies.” Evaluate this statement.

Evaluation

Chapter references

Bland, D. J., & Osterwalder, A. (2019). *Testing business ideas: A field guide for rapid experimentation*. Wiley.

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>

Haynes, L., Service, O., Goldacre, B., & Torgerson, D. (2012). *Test, learn, adapt: Developing public policy with randomised controlled trials*. Cabinet Office Behavioural Insights Team.

- Houde, S., & Hill, C. (1997). What do prototypes prototype? In M. G. Helander, T. K. Landauer, & P. V. Prabhu (Eds.), *Handbook of human-computer interaction* (2nd ed., pp. 367–381). Elsevier.
- Kohavi, R., & Thomke, S. (2017). The surprising power of online experiments. *Harvard Business Review*, 95(5), 74–82.
- McGrath, R. G., & MacMillan, I. C. (1995). Discovery-driven planning. *Harvard Business Review*, 73(4), 44–54.
- Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Crown Business.

How does an organization move from scattered initiatives that celebrate innovation to a sustained capability to produce it?

This book traces the path of innovation from the first signal to a solution adopted in practice that creates value, and analyzes the organizational foundations that make this path a settled practice: strategy, systems, culture, governance, teams and measurement, with people at the center.


It rests on three pillars: a clear, high-priority problem; motivated people whose needs are met; and a system that connects them and turns effort into repeatable results. It combines international research evidence with the Saudi context, and pairs each chapter with a tool that can be used in a real organization.

Its twenty chapters are organized in five parts: foundations; institutional foundations and people; the initiative journey; value and the ecosystem; and an application to the university innovation ecosystem.

Innovation Management

Dr. Omar A. Alsamani

ISBN 978-603-06-7961-4



9 786030 679614