

THE EMPTY CHAIR

Who's Missing from Your
Business Decisions?



VISHAL VISHWAKARMA

A customer can be right in front of you and still be missing from the decision.

“Lower the price.”

“Make it look better.”

“Give them more choices.”

Everyone has an answer.

Then the customer says, “I just need to know if it will arrive before Friday.”

A whole conversation changes with one sentence.

This book will help you find those sentences—and work out what to do with them. We’ll build useful personas, turn observations into ideas, and test those ideas before calling them improvements.

It starts with making room for the person who will live with your decision.

That’s the empty chair.



The stories are imagined. The questions are for your next decision.

I've sat in a few different chairs.

Vishal Vishwakarma

"The most important person in the meeting is the one who isn't there."

It started with questions.

How does this work?

Why didn't anyone buy it?

Could this become a business?


Each one moved me somewhere new: from writing code, to testing what makes people buy, to building my own products and small gadgets on my desk.

Most lessons came from what didn't work. A feature nobody used. A test that proved me wrong.

So I learned to take one step, check it, then take the next.

This book walks the same way.

And it starts with the chair I'm still learning to sit in: the customer's.



Where do we go from here?

We'll follow one buyer's question, then carry the way of thinking into decisions of our own.

- 01 Meet the person** / 5
Understand the situation behind a choice.
- 02 Connect the clues** / 13
Hear the request. Investigate the need.
- 03 Choose the problem** / 23
Choose what deserves attention first.
- 04 Make a prediction** / 33
Say what might help and why.
- 05 Plan a fair test** / 43
Make the comparison useful and fair.
- 06 Read the result** / 53
Understand what the evidence allows.
- 07 Make the next decision** / 63
Act, keep the learning and ask again.

Start with the person who needs the gift by Friday.



Part 01

Meet the person


Who needs it by Friday?



Imagine buying a birthday gift for Saturday. You like it. You can afford it. You still hesitate.

You need it by Friday. The page tells you what the gift is made of, how popular it is and which colours are available. It doesn't tell you whether it'll arrive in time.

The missing answer is small. What depends on it isn't.



Who's in the chair now?

The person trying to make an occasion go well.

A birthday isn't an age group.



You could describe her by age, income and job title. None of it explains why she's hunting the page for an arrival date.


Her situation does. A birthday is coming, and the gift has to be there.

She isn't really buying an object. She's trying to give someone a good birthday, on time, with something that feels right.

That's what a useful persona captures: what someone is trying to do, what matters to them and what's in the way.



Try this Finish the sentence: "This person is trying to ____, but ____ is making it difficult."



Not an age group. An occasion.

Four things worth knowing.



Keep the description small enough to use. Our first conversation gives us four notes.

Look for	Our birthday buyer
Situation	A birthday on Saturday.
Goal	A suitable gift, in time.
Obstacle	She can't find a reliable arrival date.
Evidence	"Will it arrive by Friday?"

It's one conversation. Treat it as a clue until you hear it again.



SITUATION

GOAL

OBSTACLE

EVIDENCE


The same person. Another day.



A month later, she buys something for herself. No deadline. She's happy to wait for the colour she wants.

She hasn't become a different person. The purchase has a different job.

The chair holds a person, not a label.



Is this the right one?

Someone else is upgrading a laptop that still works. Waiting a few days is fine. Their question isn't "When?" It's "Is this the right one?"

Both buyers want reassurance, just about different things.

Situation	What would help
A gift for a fixed occasion	A reliable arrival date before ordering
A planned laptop upgrade	Clear compatibility and comparisons

And if that laptop had died the night before a deadline? Suddenly delivery matters most. The product doesn't decide the need. The stakes do.



Try this Name two situations your product serves. Where do their needs split?



Not "when?" but "which one?"

One sentence is a clue.



“Can you guarantee the date?”

“It’s for a birthday.”


“I need it before the weekend.”

Three sentences. Possibly one concern.

Collect the words, where they came from and what was happening around them. Keep the examples that don’t fit too. A pattern that only survives when you ignore people isn’t a pattern.

What would make you change this description?

→ *Write the exact words down. Your memory will tidy them up.*



Give the persona a job.



In a meeting, someone suggests a larger product photograph.

Bring our buyer into the room. Would a bigger photo answer her delivery question?

Maybe the photo matters too. The persona won't approve or reject the idea for you. It makes you ask who the idea helps, and which obstacle it removes.



Try this "For someone trying to ____, we should consider ____ because ____."



BIGGER PHOTO

VS



ARRIVAL DATE

Keep a space for what you don't know.




Our description has a deadline, a goal and an obstacle. It also has gaps.

How often does this happen? Can the shop keep the date? Would a clearer date change anything?

Write those questions beside the persona. Now it guides learning, not just meetings.

Next, we'll follow her words without assuming they mean what we first think.



Part 02

Connect the clues


Faster, or more certain?

“Can you deliver faster?”

The shop starts talking about express delivery. Before anyone changes anything, someone asks, “When do you need it?”

“By Friday. It’s a birthday gift.”

The standard service arrives on Thursday. She just hadn’t found that.



Who’s in the chair now?

The person waiting for an answer.

She didn't need a faster van.



She needed an answer.

The request asked for speed. The follow-up found a deadline.



An answer before a faster van.

Ask one question further.



“Can it be faster?” can mean a birthday, an urgent replacement, or simply “I don’t know how long this takes.”

Ask, “When do you need it?” or “What happens if it arrives later?”


Then let the answer change your mind. A follow-up question that only confirms what you already believe isn’t really a question.

YOUR STORY GOES HERE

A time a client or teammate asked for one thing, and one more question showed they needed something else. 2–3 sentences, starting “I once...”



Try this Choose one request. Ask one open question that could reveal a different need.



Then she sees the delivery charge.

Our buyer reaches the total. "That's expensive," she says.

Someone suggests a discount.

Wait. Was the gift beyond her budget? Did she expect free delivery? Is she comparing it with something that arrives sooner?

Same words. Very different decisions.

→ Notice how fast "discount" shows up in a meeting.
"Why?" usually comes later.

Gift	₹1,299
Delivery	₹249
TOTAL	₹1,548

Three possible meanings.

A price comment tells you something bothered her. It doesn't tell you why.

What might be happening

What you could ask

It's beyond her budget


"What price range did you have in mind?"

She can't see why it costs more

"What are you comparing it with?"

Delivery changed the total

"Was the total what you expected?"



A clue needs somewhere to go.



Follow the words all the way to a decision.

Words: "Will it arrive by Friday?"



Need: Have the gift before the birthday.



Obstacle: The arrival date is hard to find.



Decision: What to explain before someone orders.



Try this Do this for one comment you've heard. Mark the step that's still a guess.


Can you keep the promise?

A date beside the button sounds helpful. But what if the shop can't deliver by that date?

Then the problem lives behind the screen too: stock, packing, dispatch, couriers.

Clearer words can't fix a broken promise. They only help more people rely on it.

What has to be true before you make a promise clearer?



Can the van keep it?

On the screen. Behind the screen.



What you find	Where it belongs
The service works, but people can't find the information	Make it easier to find and understand
The service can't reliably meet the need	Improve the service, or offer an honest alternative
Both	Fix both. Promise only what you can deliver.

The shop checks recent orders. Most arrive when promised. So it sets realistic arrival estimates and keeps watching for missed dates.


The problem is on the screen.



Try this Who needs to act on what you found? A writer, a designer, operations, or all three?



ON THE SCREEN



BEHIND THE SCREEN


Don't make the clue do too much.



One question suggests a problem. It doesn't prove every buyer has it.

Keep three things apart: what the customer said, what you think it means and what you want to change. Then you can rethink one without losing the others.

What do you know? What are you adding?



Keep them apart.

A better question leaves you with a choice.




We now have three possible problems: uncertain delivery, a surprise charge and limited colour choice.

All three could matter. We can't fix everything at once.

The empty chair helped us hear more clearly. Now it has to help us choose.

 **Try this** List two problems your clues support. Carry them into the next part.



Next: Everyone has an idea.

Part 03

Choose the problem


Everyone has an idea.

“Add more colours.”

“Offer a discount.”

“Make delivery clearer.”

The meeting fills with ideas. Meanwhile, our buyer still doesn't know if the gift will arrive.



Who's in the chair now?

The person who may leave while we debate.

How many people is “everyone”?




“Everyone is asking about delivery.”


You check the support inbox. Eight of twelve conversations mention arrival dates. A good clue.

But that’s only the people who wrote in. What about those who ordered without asking? Or left without a word?

Who’s missing from this evidence?



THE INBOX



THE WIDER AUDIENCE

Who had a chance to be heard?

Support hears the people who ask. Your analytics shows people searching for delivery details. Some people simply vanish at checkout.


Each source sees one part of the customer. None of them sees the whole person.

YOUR STORY GOES HERE

A time the loudest voice (a stakeholder, a complaint, one big client) set the priority, and what the wider evidence showed. 2–3 sentences.



Try this Replace “everyone” with “we heard this from ____, among ____.”



The loudest voice is not the whole room.

Four questions before a favourite.



You don't need a scoring model. You need four honest questions.

Ask	It helps you weigh
Who faces this?	How far the problem reaches
What happens because of it?	Whether it blocks a goal or just annoys
What makes us think it's real?	How strong the evidence is
What would a first step take?	The work, cost and people involved

→ *My favourite is the third question. It ends a lot of arguments.*



A request isn't an instruction.



One buyer wants another colour. Worth exploring, but it means new stock, new photos and more to maintain.

Several buyers can't find an arrival date that already exists. Making it clearer is a smaller first step.

That's a good reason to start with delivery information. It isn't proof it'll work.

Choose a next step you can explain.



*New stock, photos,
pages to maintain*



*Already exists.
Just hard to find.*

What would the wrong choice cost?



The cost isn't just the time spent building. It's the problem you left untouched.

A new colour may delight someone. But if a confusing total is pushing buyers away, the colour won't bring them back.

What are you choosing to leave for later?



What does your choice delay?

A bigger decision uses the same questions.




Now imagine the shop wants to launch a next-day gift range.

That's not a label on a page. It changes stock, packing, delivery costs and the promise people pay for.

Who needs it? How often? Can the shop deliver it profitably? What evidence would justify the commitment?

The stakes get bigger. The first question stays the same: who's in the chair?



Bigger stakes. Same first question.


Investigating is a decision too.



Not sure about colour demand? Talk to buyers and check what they searched for.

Not sure delivery is reliable? Look at recent orders.

You don't have to choose between building now and doing nothing. Sometimes the best first step is a smaller question.



Check recent orders first.

Write down why this goes first.



A good decision note can be two sentences:

“We’ll explore clearer arrival information first because buyers keep asking and the information already exists. We still need to check whether moving it helps.”

Anyone joining later can see the reasoning. When new evidence arrives, you change the decision instead of arguing about taste.

Your reason should outlive the meeting.



Choose the problem. Then make a prediction.



The shop picks arrival information first. Colour demand goes on the learning list.

That gives focus. Now we need to say what we expect the change to do.

A favourite says, "I like this." A prediction says, "Here's what will happen, and here's how we'll know."



Try this "We'll explore ____ first because _____. We still need to know _____."



I like this.

Here's what will happen.

Next: What do you expect to change?

Part 04


Make a prediction

What do you expect to change?

“Put the date beside the order button.”

That's an instruction. It isn't a reason yet.

What should become easier for her? And what should happen because of it?



Who's in the chair now?

The person who will live with our idea.

Because is doing important work.



We believe a clearer date will reduce doubt before ordering. That's why we expect more people to buy.

Keep the idea and its reason together. Without the reason, you can't learn from the result.

If it doesn't work, you'll want to know why. Was the information already clear? Was the date the problem? Did people simply not notice it?

What makes you think this will help?

→ *If you can't finish the "because", you're not ready to test.*



THE IDEA

THE REASON

An idea without a reason can't teach you.

Say it so someone can challenge it.





“This will improve the experience” hides everything.

Try this instead:

“People buying for a fixed occasion struggle to find an arrival date. We believe showing it beside the order button will help them decide and buy.”

Now someone can push back. Is the evidence strong? Is the date reliable? What counts as buying?

A clear prediction invites a useful argument.



A prediction you can reuse.

The same shape works for a product, a message or a service.

Complete the thought	Our example
For someone trying to...	Have a gift before a birthday
We noticed...	Questions about timing; trouble finding the date
We believe this change...	Showing the date beside the order button
Will help them...	Decide whether it'll arrive in time
And we'll look for...	More people completing a purchase



Try this Fill it in for one idea, in words you'd use out loud.

For...

We noticed...

We believe...

Will help...

We'll look for...


What would “better” look like?

More people click the button. The team celebrates. Sales don't move.

YOUR STORY GOES HERE

A test or launch where clicks, sign-ups or engagement went up but the real outcome didn't. 2-3 sentences.

A click is a step. What is it a step towards?



Follow the goal to a useful measure.



For the gift shop, the measure is simple: of the people who saw each version, how many bought?

For a help page, it might be how many people solved their problem. For a new product, how many reached their first useful task.

Measure success, not activity.



Success, not activity.

Agree what counts.


Does a purchase mean an order placed? A payment taken? What about an order that's later cancelled?

How long does someone get to decide? The shop picks one week from their first visit, and counts each person once.

Write this down before you see any results. Otherwise, "success" quietly drifts towards whichever version looks better.

Could two people calculate your measure and get the same answer?

→ *Boring step. Skip it, and the argument comes later, louder.*




Same definition. Same answer.

What must not get worse?

A clearer promise could bring more orders, and more missed deadlines if delivery can't keep up.

So watch what must not slip: checkout errors, cancellations, late gifts.

The main measure gives you direction. These keep you honest.



What would be worth the work?




Moving a line of text is cheap. A small gain is worth it.

A new delivery operation is expensive. It needs a much bigger gain.

Decide what “worth it” means before a promising result makes every cost look smaller.

What’s the smallest improvement that would change your decision?



Give the idea a fair chance.



We know the change, the reason, the measure and what must not get worse.

That's enough to plan a test. It's not enough to call it an improvement.



Try this Write down your change, your reason, your main measure and one thing that must not get worse.



Ready for a fair start.

Part 05


Plan a fair test

Two experiences. One question.

Some visitors see the arrival date where it is today. Others see it beside the order button.

Does the new spot help people decide and buy?

Both versions keep the same honest promise.



Who's in the chair now?

*The person choosing between our words
and their needs.*


Monday isn't Friday.

The old page ran on Monday. The new page ran on Friday.

Friday had a birthday promotion, fresh stock and twice the visitors. The new page sold more.

So what caused it?

Run both versions at the same time.



Different days. Different conditions.


Let chance choose the version.



If your keenest buyers all see the new page, it'll look great before it's done anything.

So let chance decide who sees which version. Random assignment makes the two groups comparable, without anyone hand-picking.

It won't make the groups identical. It makes the comparison fair.



Let chance pick, not people.


Keep the person in one experience.



She opens the page on her phone, then comes back later on her laptop. A different promise each time would confuse her.

Keep returning visitors in the same version wherever you can.

Your test is happening inside someone's day. Don't make it another obstacle.



Same person. Same version.

Give both ideas the same conditions.



Change one thing, and you'll know what helped. Change five, and you won't.

Keep the same

Change only

Price, stock and delivery

Where the arrival date appears

What counts as a purchase


The page location, not the promise

Who can take part, and when

Which version each person sees



Try this What could accidentally give one version an advantage?



One thing moves. Everything else stays.

Check the door before inviting people in.


Try both versions on a phone and a computer. Place a test order. Check the date, the button and what your tools record.

A missing purchase record can make a good page look weak. A broken checkout can sink a good idea.

YOUR STORY GOES HERE

A test that was broken by tracking, a bug or a setup mistake, and how you found out. 2–3 sentences.

Can you trust what this test is recording?




Don't ask the first few visitors to speak for everyone.



The first visitors may come from a promotion. The weekend brings a different crowd. Some people take days to come back and buy.

So plan for enough people and enough time before you start. A testing tool or an analyst can work out how many you need from your usual sales and the size of change you care about.

The toffees in the next part show why.



Write the finish before the start.




Decide three things up front: how many people, which days and how long buyers get to decide.

Then set a last day. If the answer still isn't clear by then, say so. Don't keep moving the finish line until your favourite wins.



Try this What makes your test complete? And what happens if it never gets there?

→ *Starting a test is easy. Stopping when you said you would is the hard part.*



You can stop to protect people.




If checkout breaks, stop. If the date becomes false, stop. If tracking fails, pause and find out why.

Agree on other red lines in advance: what triggers a pause, who decides and how the normal page comes back.

Stopping to protect customers doesn't mean the idea lost.



Try this Write one condition that would make you pause immediately, and who would do it.




Protect people first.

Now let the test teach you.



The test starts. Both versions work. Visitors are split fairly. The finish line and red lines are written down.

Every day brings a little more evidence. That doesn't mean every day needs a winner.



Next: One toffee changes the picture.

Part 06

Read the result


One toffee changes the picture.

A manager once taught me sample size with toffees. I've never forgotten it.


Picture two groups of ten visitors. Each purchase earns a wrapped toffee. A earns five. B earns six.

A: five purchases

B: six purchases



Filled toffees mark purchases. Outlines mark visitors who didn't buy.



Who's in the chair now?
The person we can't understand from one lucky purchase.

One extra purchase is the whole lead.



B is ahead. Then one of A's visitors comes back and buys.

The lead is gone.

That doesn't make B bad. It shows how easily a small pile changes shape.

Would one more purchase change your story?

A: six

B: six



One more purchase. No lead.

More people. Less sway from one person.

Now imagine hundreds of people in each group. One extra toffee barely moves the pile.

That's what sample size does: it helps you tell a real difference from ordinary ups and downs.

Small improvements need more people to spot than big ones. There's no magic number that fits every test.

YOUR STORY GOES HERE


A test with an early lead that faded, or one you were sure would win but lost. 2–3 sentences.



Try this Is your decision resting on a pattern, or a handful of toffees?



TEN VISITORS



HUNDREDS

A bigger pile of the wrong evidence is still wrong.



You could collect thousands of visits and miss every mobile purchase. Or send the new page only to people who were already going to buy.

More data won't fix that.

Enough evidence and trustworthy evidence are different questions. Check who took part and what your tools recorded before asking who won.



How much could the answer change?




Ask how long a journey takes.

“About half an hour” sounds certain.

“Usually twenty to forty minutes” is more useful.

Test results deserve the second kind of answer. Not just “B is 8% better,” but a range: how much better, or worse, it could really be.


Ask for the range, not just the headline.



What does the range allow?

You don't need to calculate the range. You do need to know what it means for your decision.

If the range shows...	Ask...
Both gain and harm	Do we know enough to choose yet?
Only gains, but some are tiny	Would the smallest gain justify the work?
Only gains big enough to matter	Were the data and customer checks sound?



A badge can't make the decision.



Your tool says “Winner.” Someone cheers.

That answers one question: is there likely a real difference? It doesn't answer the rest. What will it cost to maintain? Could it hurt the service? Does it help the people you meant to help?

Ask your tool or analyst to explain the result in plain words. You should be able to say what you're relying on.

A real difference can still be too small to matter.

→ A green badge feels like permission. It isn't.



More orders. A better decision?




The test ends as planned. B has more buyers than A.

The team celebrates. Then someone asks about late deliveries.

A gift that sells but misses its birthday hasn't done its job.

Look at what happened after the purchase, not just the purchase.



Sold. But a day late.

Read the whole result.



Before you decide, put it all together.


Check	Question to answer
Difference	How much better or worse did it look?
Uncertainty	What does the range still allow?
Value	Is the smallest likely gain worth the cost?
Experience	Did errors, cancellations or late orders rise?
Trust	Did the split and the tracking work?

What can you honestly say now?

For the shop, the answer is good. The gain is useful, the records are sound and deliveries stayed on time. It decides to roll out the clearer page.

Not every test ends this way. If it's unclear, say so. If people were harmed, fix it. If the tracking failed, make no claim at all.

 **Try this** Finish this: "We learned _____. This supports _____. It doesn't yet prove _____."



Next: A result needs a next move.

Part 07

Make the next decision


A result needs a next move.



The meeting asks, "Did we win?"

A better question: "What does the evidence let us do?"

A test isn't finished when someone reads the report. It's finished when it changes a decision.



Who's in the chair now?

The person waiting for us to do something useful.

Keep. Change. Retire. Learn more.



Keep it when the gain is useful, the records are sound and customers are fine. Roll it out gradually.

Change it when the idea still has promise but the execution let it down. Fix it and test again.


Retire it when strong evidence says it isn't worth it.

Learn more when the answer is unclear.

And if it caused harm, switch back first. Then investigate.



Try this Which one does your evidence support? Who owns it?



Unclear is an answer.

“We can’t tell yet” feels disappointing. “It doesn’t work” sounds decisive.

They’re not the same. If the range still allows both gain and harm, the question isn’t settled.


Don’t buy certainty with stronger words.

YOUR STORY GOES HERE

A result you had to report as unclear: how people reacted and what you decided next. 2–3 sentences.

What would you need to learn before your decision changes?

→ “Unclear” might be the most honest word in testing.




Roll out without forgetting the promise.



The shop rolls out carefully. The gain is worth the cost, and deliveries stayed on time.

The team keeps watching purchases, missed deadlines and customer questions as the new page becomes normal.

A test tells you about the people and conditions it saw. Stock, seasons and shoppers change.



Keep watching after launch.


Who did it help?

Someone says the new page must work best for birthday buyers.

It's a fair question. The overall result can't answer it.


If it matters, plan for it. Know who's buying for a deadline before they see either version, then compare the groups directly. One clear result and one unclear result don't prove the groups are different.

Keep a good question separate from a proven answer.



DEADLINE

VS



NO DEADLINE

Same shop. Different help.


Now personalisation earns its place. You've learned enough to ask whether different situations need different help.

A buyer who picks an occasion and a needed-by date could see delivery options first. Someone comparing materials could see a clearer comparison.

Use a real signal. Keep a useful default for everyone else. And never hide essential information from anyone.



Try this Name the signal, the different help, the default and the result you'd check.



Let the next question earn its place.




Slice results into enough groups and something will look lucky. Treat surprises as leads, then test them properly.

The same method works for bigger choices: a new service, a new price, a new product. Just make sure the test matches the question.

What would justify the extra complexity?

→ *Look hard enough and the data will agree with you.
That's the problem.*



*lucky
slice?*

Keep the learning, not just the winner.



The shop has made its decision. Now keep the learning in a short record anyone can use.

Step	What to keep
Find	The person, the situation and the evidence
Learn	The need, the barrier and the open questions
Predict	The change, the reason and the measure
Test	Who took part, the fair comparison and the finish line
Decide	The result, its limits, the next step and its owner

A test that loses but teaches you something is still worth keeping.



Try this Use these five steps on your next decision.

Leave room for the next person.




This book began with a customer missing from a decision. It ends with a habit for the next one.

Before a feature, a message, a service or a big commitment, ask three things. Who does it help? What are they trying to do? Why do you think it'll help?

You'll still make assumptions. The real work is noticing them, checking them and changing your mind when the evidence says so.

Who belongs in the chair at your next meeting?



Who sits here next?

Thursday.



Picture our birthday buyer again. The clearer page is in front of her. Beside the order button: *Thursday*.


She orders. On Thursday, the gift arrives. On Saturday, it's opened.

The test helped the shop choose a clearer page. The service kept the promise. Together, they helped her do what she came to do.

At the next meeting, everyone may still have an answer.

This time, someone asks who will live with it.

That's how we fill the empty chair.



Names for ideas you have already met.

Persona

A working description of a recurring customer situation, grounded in evidence. Its purpose is to guide a decision, not decorate a presentation.

Hypothesis

The prediction connecting an observed problem, a proposed change, a reason and an outcome you can check.

KPI

Key performance indicator. Your primary KPI is the main outcome you chose before evaluating the change. Define exactly what is counted and over what period.

Guardrails

Measures of what must not get worse while you pursue the main outcome. Give them clear definitions, owners and action thresholds.

Return to the story whenever a label makes the idea harder to see.

Names for a fair test.

A/B test

A randomised comparison of two experiences, usually the current version and a proposed change, evaluated against outcomes chosen in advance.

Eligibility and assignment

Eligibility states who can take part before assignment. Random assignment decides which experience they receive. Keep outcomes for the assigned groups without cherry-picking afterwards.

Denominator

The people or opportunities counted underneath a rate. More purchases can reflect more visitors, rather than a better experience. Compare outcomes relative to comparable opportunities.

Sample ratio mismatch

An unexpected difference between the planned and observed group split, beyond ordinary chance. It can signal allocation or tracking problems and needs investigation before trusting results.

Test the setup as well as the idea.

Names for what you can and cannot tell.

Sample size

How many units, such as people, are included. The sample needed depends on the outcome, baseline rate, effect worth detecting, allocation and statistical method.

Statistical power

How likely the planned test is to detect a specified effect if that effect exists, under the assumed model. It is not the probability that your idea is right.

Confidence interval (the range)

A range expressing uncertainty using a chosen method. A 95% method would cover the true effect in about 95% of repeated studies under its assumptions, not promise certainty for one decision.

Statistical significance

Whether the result crosses the planned evidence threshold against a stated null hypothesis. It does not establish that a gain is large enough to be useful.

Ask an analyst or your platform to explain its method in language you can use.

Names for the next move.

Kill criteria (red lines)

Pre-agreed conditions for stopping or pausing to protect customers or test validity. State the trigger, owner and response. A safety stop is not automatically a statistical loss.

Fixed-horizon and sequential tests

A fixed-horizon test uses a planned final analysis. A valid sequential method supports repeated statistical decisions while controlling error. Choose the approach before testing.

Segment analysis

Studying effects within groups. Use meaningful signals known before assignment. A significant result in one group and a non-significant one in another do not establish different effects; compare effects directly.

Personalisation

Tailoring an experience to a relevant situation or signal. Keep a useful default. Searching many groups creates false-positive risk; confirm unexpected findings with a new planned test.

The purpose of the terms is a better decision, not a more complicated conversation.

What sits behind these pages?

The gift shop, its customers and its outcome are imagined teaching examples. The toffee explanation draws on an example a manager shared with the author. These stories show a process, not a promise of results.

The author's work and projects shape the questions. These sources support the experimentation and uncertainty concepts.

GOV.UK

A/B testing: comparative studies

<https://www.gov.uk/guidance/ab-testing-comparative-studies>

Larsen and colleagues

Statistical Challenges in Online Controlled Experiments (2023 revision)

<https://arxiv.org/html/2212.11366v5>

NIST/SEMATECH

e-Handbook: Confidence intervals

<https://itl.nist.gov/div898/handbook/prc/section2/prc241.htm>

First reader draft / October 2026. Author photograph to be supplied.

What made you think?

Page / idea / question

Use page numbers to connect your notes to the draft.

What would you change?

Page / idea / question

Use page numbers to connect your notes to the draft.