

FROM THE AUTHOR OF THE EMPTY CHAIR

The Flip-Up Problem



How one small customer complaint can fix
your product, your page and your sales.

VISHAL VISHWAKARMA

Every complaint is a clue.



My first book, *The Empty Chair*, was about making room for the customer before you decide. This one is about what to do once they're in the room.

Most growth advice starts with tools and dashboards. This book starts with a person: a rider who loves his helmet but can't take a call in it.

Follow him through seven short parts and you'll meet the words people use in meetings: personas, A/B tests, sample size, confidence, personalisation. Each one comes with a story, not a definition.

Every part ends with one question to ask in your next meeting, and one prompt to try with AI on your own data.

You don't need to be technical. You need to be curious about one customer at a time.

Vishal Vishwakarma



Every part, the same six steps.



- ① **The story** One brand. One real problem.
- ② **What it is** The idea in plain words, no jargon.
- ③ **How it helped** Product, experience and sales, side by side.
- ④ **How you do it** Three to four simple steps for your business.
- ⑤ **Ask this in your next meeting** One sharp question that changes the room.
- ⑥ **Ask AI** A ready prompt to dig into your own data.



Start anywhere. Each part stands on its own, and they build on each other in order.

Big words. Simple ideas.



Word	What it means	Example	Part
Persona	The real person behind the order.	The rider who needs to take calls.	01
Find · Learn · Decide	Where someone is right now.	Just looking, curious, or ready to buy.	02
Reading the room	Letting data you already have pick where to start.	8 in 10 visit on a phone.	03
A/B test	Two versions at the same time. Keep the winner.	Flip-up first vs the usual page.	04
Sample size	Enough people that one can't flip the answer.	1,000 visitors, not 10.	05
Confidence	How sure you are the gap is real, not luck.	A lead of 200, not 4.	05
Personalisation	Same shop, different doors.	Callers see flip-ups first.	06

The promise: no jargon without a story. Every term comes with a real example.

7 stories. 7 better questions.



01	Personas	6
	The real person behind the order	
02	Find · Learn · Decide	14
	Where your customer is right now	
03	Read the Room	20
	Let the data you already have show where to start	
04	Test, Don't Guess	26
	One idea at a time, proved with data	
05	Enough to Trust	33
	Sample size and confidence, made simple	
06	Same Shop, Different Doors	41
	The right thing for the right person	
07	The Whole Loop	47
	One brand, start to finish	

01



Personas

The real person behind the order

Two helmet companies. One question.



Company A

Already sells thousands of helmets. Has a hero product and loyal customers.

Asks: How do we grow beyond the customers we already have?


Company B

A new brand, about to launch its very first helmet.

Asks: How do we get in when the market is already crowded?

Two different businesses. You'd expect two different answers.

Both answers are hiding in the same place: one rider's day.



Two businesses. One rider's day.

Meet the rider.

He isn't an age group or a pin code. He's a person with a routine.

Time	His day
8:30	Rides 40 minutes to work
9:05	Phone rings at a red light
9:15	Thirsty in slow traffic
19:00	Rides back home, same again

"I love my helmet. The quality, the look, all of it. But the front doesn't open. To sip water or take a call, the whole thing has to come off."

He loves quality, look and fit.

He's missing a front that opens.



Calls and water, on the move.

One rider. Two opportunities.



Call him the On-the-Go Rider. He wants to talk and drink without taking his helmet off.

Company A

Already has a flip-up? Show it to this rider first: in ads, on the homepage, on the product page.

No flip-up yet? Add one as a variant, and win back a customer type they were quietly losing.

Company B

Launch for him. “The helmet you can talk and drink in.”

Enter through a gap, not a price war.

Same person. Two businesses. Both grow.



Talk and drink, helmet on.

So, what is a persona?



A persona is one type of customer: what they do, what they want, and what gets in their way.

What reports say	What a persona says
Male	Rides 40 minutes to work
25 to 34 years	Takes calls on the move
Delhi NCR	Gets thirsty in traffic
Shops on mobile	Will pay more for an open front

Demographics tell you who.
A persona tells you why.



SITUATION

GOAL

OBSTACLE

EVIDENCE


From many riders to a few personas.



Listen to many riders first. Group what repeats. Then give each group a name based on what they do, not who they are.

What riders say	Persona and need	Our helmet
"Can't take calls." "Need to drink water."	On-the-Go Caller An open front	Gap
"Sun in my eyes." "Glare at 5 pm."	Sun-Glare Rider A built-in sun visor	Partly
"Hot on long rides." "Neck hurts after 3 hours."	Long-Haul Tourer Light and airy	Fits

Now you can see exactly where your product fits, and where it lets people down.



What changed for the helmet brand.



Company A added a flip-up variant of its hero helmet. Nothing else about the brand changed.

**PRODUCT**

A flip-up version of the helmet riders already loved.


**EXPERIENCE**

Callers could finally find a helmet built for their day.

**SALES**

A customer type they were losing started buying again.

Conversion was never the goal. It came from solving a real problem.



Growth from a gap, not a price war.

Build your first persona.



- 1 **Listen.** Reviews, returns, support calls, staff on the floor.
- 2 **Group.** Three to five groups, not one.
- 3 **Name.** By behaviour, not by age.
- 4 **Check.** Where does our product fit each one?

ASK THIS IN YOUR NEXT MEETING


“If we spent one day as each of our customers, where would our product fit in, and where would it let them down?”



ASK AI

“Here are 200 reviews of our helmet. Create four personas based on how these riders live and ride. For each one, tell us how well our helmet fits their need and where it falls short.”

AI gives you a first guess. Check it with five real customers before you build on it.



Where your customer is right now

Three people walk in.

Same helmet shop, same afternoon. Three very different moments, in the store and on the website.

What they say	In store	Online
"Just looking around."	Wanders along the helmet wall	Lands on the homepage, scrolls
"Can I really take calls in this?"	Picks one up, asks the staff	Reads the product page, watches the video
"Is it in M? Can I pay by UPI?"	Walks to the counter	Adds to cart

A good shopkeeper reads all three: greets the first, explains to the second, bills the third fast.

Does your website know who's who?



Same shop. Three moments.


Find. Learn. Decide.



Three moments. Each one needs something different from you.

Moment	In store / Online	Your one goal
Find Just looking	Helmet wall, window / Homepage, landing pages	Help them see the right options.
Learn Curious	Staff demo, trying one on / Product and feature pages	Answer their questions.
Decide Ready	Counter, size and price / Cart, checkout	Make buying easy.

Know the moment. Give it one goal.




Right person. Wrong moment.

Same three people. This time, everyone gets the wrong thing.

Moment	What happened	What they needed
Find	Pushed to buy at the door. A “Checkout now!” pop-up.	Options to look at.
Learn	“Bestseller!” but no answers. No feature details, just Buy.	Answers to their questions.
Decide	Long queue, card machine down. Forced sign-up, surprise fee.	A fast, easy way to pay.

The right message at the wrong moment is still the wrong message.



He was just looking.

What changed for the helmet brand.



Our caller spends most of his time in **Learn**. He wants proof he can really take a call. So the brand gave the flip-up product page one job: answer that question.

**PRODUCT**

A short clip of a rider taking a call with the front flipped up.


**EXPERIENCE**

The question “Can I talk in it?” answered before he had to ask.

**SALES**

Fewer riders leaving the product page to search somewhere else.

answered



“Can I talk in it?”

Sort every spot into a moment.



- 1 **List every spot.** Store: window, wall, counter. Online: every page.
- 2 **Sort.** Is this spot for Find, Learn or Decide?
- 3 **Give each one goal.** Window: show the range. Checkout: make paying easy.
- 4 **Check.** Does each spot do its one job, and nothing else?

ASK THIS IN YOUR NEXT MEETING

“Who is this page for: Find, Learn or Decide? And what’s its one job?”



ASK AI

“Here are our website pages and store spots: [LIST]. Sort each one into Find, Learn or Decide, and suggest one clear goal for each.”

03



Read the Room

Let the data you already have show where to start

Everyone loved the new design.



The helmet brand reviews its new flip-up page on the big screen in the meeting room. Big photo, clean layout. Everyone nods.

What the team saw

A big, beautiful photo

“Looks great.”


What customers saw

A photo that filled the whole phone screen

“Where’s the button?”

Eight in ten customers visit on a phone. The data had it all along.

**They checked their room.
Not their customers’ room.**



2 IN 10



8 IN 10

Let the data pick where to start.




Before you design, decide or test, look at what your data already says about who shows up, and how.

If most visitors are on...	Then...
Phones , like our helmet shop: 8 in 10	Start mobile. Design, review and test on a phone first.
Desktops , like a business software firm	Start desktop. Then check mobile still works.

Online, it's in the device report of your analytics. In store, watch which door people use, and when. In the review meeting, look at the page on the screen most customers use.

Start where most of your customers already are.



Start where most of them are.

What else is your data saying?



Clue	In store	Online
Busiest hours	Rush is 6 to 9 pm. Add staff then.	Lunch is the peak. Launch at noon.
First thing seen	The side door is busiest. Dress it.	Most land on product pages. Fix those.
Where they're from	Office crowd at 9 am.	Instagram visitors: lead with looks.
What sells	Bestsellers at eye level.	Bestsellers at the top of the homepage.

None of this needs new research. It's sitting in reports you already have.



HOURS



FIRST SEEN



FROM



SELLS

What changed for the helmet brand.




The team stopped reviewing on the meeting-room screen. Every page was checked on a phone first, the way riders see it.

 **PRODUCT**
The flip-up's open front shown in the first screen on mobile, not the third.

 **EXPERIENCE**
The "Add to cart" button visible without scrolling.

 **SALES**
A page that finally worked for the 8 in 10 riders on phones.



no scrolling



Read your room in four steps.



- 1 **Pull four numbers.** Devices, busiest hours, first page or door, top sellers.
- 2 **Find the biggest group.** Where do most of your customers show up?
- 3 **Start there.** Design, fix and review for that group first.
- 4 **Then test it.** Prove the change works. That's Part 04.

ASK THIS IN YOUR NEXT MEETING

“Are we looking at this the way most of our customers will see it?”



ASK AI

“Here is last month's data: [DEVICES], [BUSIEST HOURS], [TOP ENTRY PAGES], [TOP SELLERS]. Where are most of our customers, and what three things should we fix first for them?”

One idea at a time, proved with data




Everyone has an idea.

Sales of the flip-up are flat. The helmet brand calls a meeting. Everyone around the table has an answer.

Who	Their idea
Sales head	"Run a 20% discount."
Designer	"A new homepage video."
Store manager	"Move the flip-ups to the window."
Marketing	"Show the flip-up first to riders who searched for a calling helmet."

Every idea sounds right. Every idea costs time and money.

Everyone has an opinion. Nobody has proof.



Test one idea at a time.




The brand could try all four ideas this month. Sales might even go up 12%.

But which one did it? Nobody knows. Maybe the discount gave away money the other three earned.

So they try one idea at a time: change one thing, keep everything else the same, and watch what happens.

Change one thing. Watch one number.



One idea. Then the next.

Write it down first.



Before changing anything, the team fills in a small test card. Everything from Parts 01 to 03 fits on it.

Test card No. 1	Flip-up first
The idea	Show the flip-up first to riders who searched “calling helmet”.
Who it’s for	The On-the-Go Caller (Part 01)
Their moment	Learn: comparing helmets (Part 02)
Where	Mobile first: 8 in 10 visits (Part 03)
One number	Share of these visitors who buy within 7 days
How long	15 days, so one busy day can’t flip it (Part 05)
Compared with	The usual page, at the same time

Was it the flip-up, or the monsoon?



The team's first plan: compare the 15 days before the change with the 15 days after.


Sales went up. But the monsoon also arrived, and riders were suddenly shopping for helmets they could see out of in the rain.

Before-versus-after can't tell you which one did it.

Online, split visitors at random *at the same time*: half see A, half see B. Rain, paydays and holidays hit both halves equally.

In store, you can't split one doorway, so switch A and B on alternate days. It isn't perfect, but weather and paydays fall evenly on both.

Compare at the same time, not with last month.



What the test showed.



After 15 days, the team counts riders who searched “calling helmet” and bought within a week.


Version	Visitors	Bought	Rate
A · usual page	1,000	100	10%
B · flip-up first	1,000	130	13%

**PRODUCT**
One clear answer about what callers want to see first.

**EXPERIENCE**
Callers saw the right helmet without hunting for it.

**SALES**
B is ahead by three points.

Three points is exciting. But is it real, or luck? That's the next part.



Run your first test.

- 1 **Pick one idea.** Start with the one your data points to.
- 2 **Write the test card.** Who, where, one number, how long.
- 3 **Run A and B together.** Same days, same weather, one change.
- 4 **Check you can trust it.** Keep it, drop it, or test longer. That's Part 05.

ASK THIS IN YOUR NEXT MEETING


“Before we change anything: how will we know it worked?”



ASK AI

“We want to test [IDEA]. Help us fill in a test card: who it's for, their moment, the one number to watch, how long to run it, and what to compare it with.”

05



Enough to Trust

Sample size and confidence, made simple

Ten visitors. Two pages.



Imagine the test had only run for an hour. Ten riders saw each page.

A  **4**

B  **6**

Each filled helmet is one sale. A sold 4. B sold 6.

“B got more. B wins!”

But what if just one rider had done something different?

Now change one sale.



One rider on B was about to buy, then his phone rang. He didn't.

One rider on A was on the fence, then bought.



Now it's 5 against 5. Nothing about the pages changed. Two riders' moods did, and the whole story flipped.

If one rider can flip the story, when can you trust it?

Now count a thousand.




Same split, but this time 1,000 visitors. Move one sale from B to A.

	Before	After one sale moves
10 visitors	4 vs 6	5 vs 5. The story flips.
1,000 visitors	400 vs 600	401 vs 599. The story holds.

With a thousand people, one rider's mood barely shows.

Sample size: count enough people that one person can't change the answer.



10 VISITORS



1,000 VISITORS

Now look at the gap.




Two tests, 1,000 visitors on each page. B wins both times. But look at how much B wins by.

Test	A vs B	B's lead	Could it be luck?
One	400 vs 600	+200	Very hard to call it luck
Two	498 vs 502	+4	Could easily be luck

Confidence: how sure you can be that the gap is real. The smaller the gap, the more people you need.

Sample size asks: how many people is enough? Confidence asks: how big a difference is enough?



Is the flip-up lead real?


Back to Part 04: 100 of 1,000 bought on A (10%), 130 of 1,000 on B (13%).

- 1 **Find the gap.** $13\% - 10\% = 3$ points.
- 2 **Find the wobble.** How much luck alone could move it: $\sqrt{(10 \times 90 \div 1,000 + 13 \times 87 \div 1,000)} \approx 1.4$ points.
- 3 **Is the gap more than twice the wobble?** 3 is more than $2 \times 1.4 = 2.8$. Yes.

Yes means it passes the usual 95% check: luck alone is unlikely to explain a gap this big.

With only 100 visitors each, the wobble grows to 4.5. Same 3-point gap, not enough. Keep counting.

Don't like maths? Search "A/B test significance calculator" and type in the same four numbers.



What waiting saved.


On day three, B was ahead by eight points. Someone wanted to roll it out that afternoon.

The team waited for the full 15 days. The lead settled at three points. Smaller, but real.

 **PRODUCT**
A decision built on 2,000 riders, not a lucky Tuesday.

 **EXPERIENCE**
No flip-flopping pages while the team argued over early numbers.

 **SALES**
A real three-point lift they could plan around, not an eight-point hope.



When can you trust the data?

Before anyone says “B wins”, check three things.


- 1 **Enough people?** Could one person flip the result? Then wait.
- 2 **A big enough gap?** Could luck explain it? Then wait.
- 3 **A fair amount of time?** At least a full week, so weekdays and weekends both count.

ASK THIS IN YOUR NEXT MEETING


“Do we have enough people, and is the gap big enough to trust?”

★ ASK AI

“Version A had [VISITORS] visitors and [SALES] sales. Version B had [VISITORS] and [SALES]. Is the difference big enough to trust? If not, roughly how many more visitors do we need?”



Draw the line before the race.



Same Shop, Different Doors

The right thing for the right person

Three riders. One banner.




Remember the three personas from Part 01? They all visit the helmet shop on the same day.

Everyone sees the same thing in the window and on the homepage: **“New arrival: the Racing Helmet.”**

Rider	Needs	Sees the banner and thinks...
On-the-Go Caller	An open front	“Not for me.”
Sun-Glare Rider	A sun visor	“Not for me.”
Long-Haul Tourer	Light and airy	“Not for me.”

The shop has the right helmet for each of them. Nobody sees it.

One message for everyone speaks to no one.




Same shop. Different doors.



Same stock. Each rider gives a clue. Each gets their own way in.

Rider	The clue	Their door
Caller	Asks about calls. Searched “calling helmet”.	Flip-ups first
Sun-Glare	Squints, asks about glare. Came from a visor ad.	Sun-visor helmets first
Tourer	Talks about weekend trips. Reads the touring blog.	Light touring helmets first

Personalisation: same shop, same products, a different first thing for each kind of person.



Open one door at a time.




Open a door on a guess, and some people walk into the wrong room.

Guessing: the shop decides every city rider is a Caller. Everyone sees flip-ups first. The Sun-Glare Rider can't find a visor, and leaves.

Testing: only people who searched “calling helmet” are split in two. Half see flip-ups first, half see the usual banner. You already know how this ends: Parts 04 and 05.

A wrong door is worse than no door.





Wanted a visor. Got flip-ups.

What changed for the helmet brand.



The flip-up door stayed open for callers. Everyone else still saw the usual homepage, so nobody lost their way in.



Open your first door.



- 1 **Spot the clues.** What people search, ask, click or buy.
- 2 **Match clue to persona.** Which kind of customer is this? (Part 01)
- 3 **Choose their door.** What should they see first?
- 4 **Test before you keep it.** One door, A against B. (Parts 04 and 05)

ASK THIS IN YOUR NEXT MEETING

“What does this customer see first, and is it the right thing for them?”



ASK AI

“Here are our personas: [PERSONAS]. Here are the clues we can see: [SEARCHES, ADS, PAGES VISITED]. Match each clue to a persona, and suggest what each one should see first.”




One brand, start to finish

Six steps. One loop.



Every part of this book is one step. Put them together and they make a loop.



Each time round, you learn something new about your customers, and the next loop starts smarter.



Growth isn't a straight line. It's a loop.

One brand. One full loop.



Step	What the helmet brand did
1 · Meet the person	Reviews kept saying: can't drink water or take calls.
2 · Their moment	Callers were in Learn, reading the product page.
3 · Read the room	8 in 10 read it on a phone. Start there.
4 · Test one idea	Half saw the flip-up first. Half didn't.
5 · Trust the data	15 days later, the gap beat the wobble.
6 · Open the door	Calling-helmet searchers now see flip-ups first.

Next loop: the Sun-Glare Rider.





next loop: a sun visor

What one complaint did.



It started with one rider who couldn't take a call at a red light.



Every complaint is a door.

Start your first loop.



- 1 **Pick one complaint.** The one you hear most, or the one that costs most.
- 2 **Walk it round the loop.** Person, moment, room, test, trust, door.
- 3 **Write down what you learn.** It's where your next loop starts.

ASK THIS IN YOUR NEXT MEETING

"Where are we in the loop right now, and what's the next step?"



ASK AI

"Here's what we know: [REVIEWS, DATA, TEST RESULTS]. Walk us through the loop: who to serve next, what to test, and how we'll know it worked."

THE END. AND THE START.

Every complaint is a door. Go open one.

Start with one person. One idea. One fair test. Then go round again.

In my first book, *The Empty Chair*, I asked you to keep one chair empty in every meeting, for the customer who isn't in the room.

Now you know who sits in it: a rider at a red light, phone ringing, helmet on.



What sits behind these pages?



The helmet brand, its riders, the meetings and every number in this book are made-up teaching examples. They show a way of thinking, not a promise of results.

The wobble check in Part 05 is a simplified version of the standard test for comparing two percentages. It's good for building intuition. For real decisions, use a significance calculator or ask an analyst.

AI prompts give you a fast first draft. Check what AI suggests against real customers and real data before acting on it, and never paste in personal customer details.

Further reading

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*

<https://arxiv.org/abs/2212.11366>

NIST/SEMATECH, *e-Handbook: Confidence intervals*

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

First reader draft / October 2026.

What made you think?



Page / idea / question

What would you change?



Page / idea / question