Interviewing without bias

Stage 2, Customers and value. Lesson U2.3. How to talk to the people you want to serve so that you learn what is true, not what is polite.

> **Your task.** Ask about the last relevant choice or experience. Discovery questions remain neutral; later concept/prototype exploration is labelled as such, not misrepresented as unprompted discovery.

Why this matters

People are kind. Describe your idea and ask whether they would use it, and most will say yes, because no feels rude and the future is easy to promise. That answer feels like validation and is worth almost nothing. The interview is your main instrument in Stage 2, and a badly built instrument produces confident nonsense. This unit teaches you to build a better one.

Learning outcomes

By the end of this unit you can:

1. Write an interview guide, built from your learning goals and believed customer profile, that contains no leading, hypothetical or pitching questions and at least two questions about current spending.
2. Run an interview that surfaces specific past behaviour, current alternatives and money, time or effort already spent.
3. Record what you heard against a specific claim, with its commitment level and source details, within an hour of the interview.
4. Use an AI critique and an AI practice persona without treating either as evidence.

The lesson

Building is easy now. Knowing what to build is not.

Building something is different from establishing that people value it enough to change their behaviour. Investigate actual experiences and commitments before drawing conclusions about demand.

AI has made this risk larger, not smaller. You can now build a polished website, app or brand in a weekend. Building is no longer the hard part. **Knowing what is worth building is.**

This sprint is how you find out. Over one to two weeks you will talk to real people who have the problem you care about. The number depends on your venture type; for a consumer venture it is at least ten. Your goal is not to be told your idea is good. It is to learn what is actually true about their lives.

Why people mislead you, without meaning to

People are kind. If you describe your idea and ask whether they would use it, most will say yes, because saying no feels rude. They are not lying; they are being polite and imagining a future that will never arrive. That is why "Would you buy this?" is one of the most dangerous questions in entrepreneurship. It produces data that feels like validation and is worth almost nothing.

The fix is simple to describe and hard to practise: **ask about their life, not your idea; ask about the past, not the future; and listen much more than you talk.** (These principles are popularized in *The Mom Test* by Rob Fitzpatrick. Read it.)

What the sources agree on

The approach popularised by Rob Fitzpatrick in *The Mom Test* (2013), the best-known book-length guide to customer conversations, is worth reading in full. Its principles are taught here as general practice in Foundiro's own words, alongside the sources this programme has read directly, because they all point the same way:

- **Talk about their life, not your idea.** Bill Aulet describes primary market research as engaging real potential customers with an open mind, not selling to them (MIT 15.390, lecture 6). Steve Blank's customer discovery is inquiry, not a sales call.
- **Ask about a specific past moment, not what usually happens or what might happen.** The d.school's interview-for-empathy card makes the same point: ask about a specific occasion and keep the word "usually" out of your questions (*Design Thinking Bootleg*, the interview-for-empathy card, paraphrased). General answers ("I always...") describe what people wish they did; stories describe what they did.
- **Listen far more than you talk.** Allow silence. Ask "why?", "what happened next?", "can you show me?". The d.school card adds two habits: keep asking why, even when you think you already know the answer, and look for inconsistencies between what people say and what they do (same card, paraphrased).
- **Compliments and promises about the future are not data.** Josh Kaufman's advice: ask real potential customers rather than friends and family, ask open questions and listen more than you talk, compare what people do with what they say, and treat a real commitment such as a pre-order as far stronger than stated interest; indifference is a worse sign than criticism ("Feedback", personalmba.com).
- **Commitment signals real interest.** What people have already given up (money, time, data, reputation) tells you more than what they say they will do. In Stage 2 that commitment is about their *current* workaround, because you have not offered anything.

Five rules for every interview

1. **Talk about their world.** "Tell me about the last time you had to [task]." Not "What do you think of my app?"
2. **Go to a specific past moment.** "When did that last happen? Walk me through it." Specific stories reveal what people actually do. General statements ("I always...") reveal what they wish they did.
3. **Find the workaround and its cost.** "What do you use today? What does that cost you, in money, time, or stress?" If nobody is trying to solve the problem today, it may not hurt enough.

4. **Listen 80%, talk 20%.** Silence is your friend. Follow up with "Why?", "What happened next?", "Can you show me?"
5. **Don't pitch.** The moment you explain your solution, the honest data stops. If you must show something, do it only at the very end, and ask for a commitment rather than an opinion.

Rewriting bad questions

Before asking about the last time something happened, check that it has happened. Allow no problem, no loss, no helper and no spending as valid answers. For example: "Has a payment ever failed to match an order? If yes, tell me about the most recent time. If not, how do you check?"

Bad question	Why it fails	Better question
"Would you use an app that...?"	Hypothetical, invites politeness	"How do you handle this today?"
"Do you think this is a problem?"	Leading; they will agree	"Walk me through your most recent [task]. Did any part go differently from what you expected? If yes, what happened?"
"How much would you pay for...?"	Future and hypothetical	"What do you spend on this now, including tools, services or extra time?"
"Wouldn't it be great if...?"	Pitch disguised as a question	"What have you tried to make this easier? What happened?"
"Do you like our idea?"	Asks for a compliment	(Don't ask.) "Who else should I talk to about this?"

Money questions (at least two in every guide)

- "What do you currently pay: in money or time: to deal with this?"
- "Have you ever paid for a solution, a service, or someone's help with this? What happened?"
- "Who decides on spending for this, and how does that decision get made?"
- "During that specific occasion, how much time or money, if any, did it cost you?" Allow no cost as a valid answer; record reported costs separately from your own estimates.

Build the guide from what you believe

Value Proposition Design's Journalist exercise starts by sketching and ranking the customer profile you *believe* in, then derives the interview questions from it: ask yourself what you want to learn (The Journalist). After each interview it asks you to **review your questions** in the light of what you learned. Do the same: the customer profile you believe in (next lesson) and your claim set (first lesson of the stage) give you at most three learning goals, and every question serves one of them.

When you slip

Everyone pitches by accident at least once. Anything said after you described a solution is tagged **E2 (stated intent) at most**, unless it came with a real commitment. Note the minute you slipped in your log.

Interview purpose belongs to a claim. An empathy-style conversation ("tell me about your work") can support C1 at most. If you want evidence for C2 or C3, your questions must reach severity, frequency and spend.

During and after the interview

- Have a partner take notes. Record only with separate permission for recording.
- Write down **exact phrases** customers use. Use these words to understand the experience. Public quotation or marketing use needs separate explicit permission for that use.
- Within one hour, complete the Evidence Ledger entry while your memory is fresh, and write the time you completed it in "Entry written at"; a later entry is flagged there with the reason.

Recording evidence: claim, commitment, source

For every meaningful statement or observation, record three things.

1. Which claim is it evidence for?

Claim	What it says
C1	The problem exists for your segment
C2	It is severe or frequent enough to act on
C3	People spend money, time or effort on a workaround today
C4	They are willing to switch (later in this stage)

In this sprint you are *not* testing whether people want your solution (C5) or will pay your price (C6). You have not shown them an offer. That comes in Stages 3-5.

2. How much commitment does it show? (the E-ladder)

Level	Meaning	Example
E5	Money or repeat behaviour	"I pay a bookkeeper 500 CU a month to do this."
E4	Costly commitment	They gave real time or shared real data or documents
E3	Specific past behaviour	"Last Tuesday I spent two hours fixing order mistakes."
E2	Stated intent	"I'd definitely try something like that."
E1	Opinion or compliment	"Great idea!"
E0	AI or synthetic	Anything from an AI persona or simulation. Never evidence

3. Did you see it, or were you told? *Observed* (you watched it or saw the document) is stronger than *reported* (they told you).

Two examples make the difference clear:

- "I pay my niece 300 CU a month to manage orders" is a **reported E5 for C3**. It is strong evidence that the problem costs money today. It says **nothing** about whether this seller wants or would pay for your solution.
- A seller showing you a week of chaotic order screenshots is **observed evidence for C1 and C2**. It is research participation, not a buying commitment.

You will defend your tags in the Evidence Court.

Worked example (fictional): what changed between the first guide and the revised guide

> **Fictional teaching case.** Simulated evidence for practice.

MarketLink's first guide and its improved guide are both in the Interview Guide worksheet (your worksheets, file C, with the change log); the sprint learner workbook's worked example (section E) shows them too. Three changes did most of the work:

- "Would you use an ordering app?" became "Walk me through yesterday's orders." The question moved from their opinion of a product to their own past.
- "How much would you pay?" became "Do you pay anyone or anything to help? What does that cost you a month?" A hypothetical price became a current spend (C3).
- "What features do you want?" was **deleted**, not rewritten. Feature wishes are ideas, not evidence of a problem.

A slip, handled honestly. In interview #2, Daniel (one of MarketLink's two founders) described "an app that matches payments for you" at minute 12. The seller said "that would be brilliant, I'd pay for that". Daniel tagged everything after minute 12 as E2 at most, and the team used nothing from that part of the interview in any claim decision.

What they decided not to do. They did **not** show their ordering-app mock-up at the end of the interviews "just to get a reaction". With no offer to commit to, a reaction would be an opinion (E1-E2) and would have coloured the rest of the sprint. The mock-up waited for Stage 3.

Guided practice in the session

1. **Bad to good rewrite (30 min, pairs).** Rewrite the ten questions in the practice set in your worksheets (file C), or ten your facilitator provides. Push every rewrite towards "the last time...".
2. **Guide drafting with live critique (30 min).** Draft your first guide from your learning goals. Your facilitator demonstrates the routine of attempt, critique and compare on one team's draft.

Apply it to your venture

1. **Finish the improved guide** (30 min): 8-12 questions, at least two questions about current costs, including one about money actually spent, plus "How did you look for a fix last time?" and the two Moment-of-Truth questions ("What almost stopped you?" and "What made you go ahead?", asked about the current solution). Use the filled example in your worksheets (file C) as a model.
2. **Rehearse twice** (30 min): once with a peer, and once more with either an AI persona (practice only, never logged as evidence) or a second peer or your facilitator; the two routes are equal. Write a three-line reflection: what you would ask differently.
3. **Run your interviews** (~3 h 30 min for 4 interviews, including travel and notes). Each founder conducts some personally.
4. **Ledger within one hour of each interview** (~15 min each): claim, E-level, source details, current alternative or proposed offer, the five permission fields (research notes, recording, AI transcription, external AI analysis, public quotation: each yes, no or pending, with scope and date), item class, and "Entry written at" (the time you wrote the entry; a late entry is flagged there).
5. **Review your questions after each interview** (5 min, inside step 4): which question produced a story, which produced an opinion?

If you are creating something new

Ask what made the occasion valuable and what the person actually chose; do not lead them toward a pain they did not report.

Working with Your AI Partner

- **May do:** flag leading, hypothetical and pitch questions in *your* draft and check for two questions about current costs, including one about money actually spent; play an evasive practice interviewee and rate your questioning; transcribe and redact **consented** recordings.
- **You must check:** every suggested rewrite (accept or reject, with a reason); every transcript against the recording before any line enters the ledger.
- **Human first:** draft your first guide before any AI critique.
- **Never:** write the guide from scratch, or have any persona output logged as evidence. Persona output is E0.
- **No AI available?** A peer critiques your guide against the banned list on worksheet 1 ("would you / do you think / how much would you pay / pitch"); rehearse with a peer only.
- **Log:** an AI use note (worksheet 5 in file C) on the guide: what the critique flagged, what you changed, what you rejected.

Synthetic users are not customers

AI can now produce convincing customer personas, simulated interviews, even "market research reports" in seconds. These are useful for rehearsal and for generating hypotheses. They are never evidence. An AI persona can only recombine what is already written somewhere. It cannot tell you that *this* seller paid her niece last month, that a clinic manager has a budget code for the

problem, or that the person you expected to care doesn't.

A useful habit, drawn from *AI Snake Oil* (Narayanan & Kapoor): when a tool claims to "predict demand" or "simulate your market", ask three questions. What exactly does it predict? What evidence shows it works for cases like yours? How would you know if it were wrong? Then treat its output as a hypothesis to test with real people.

The founders who win in an AI-saturated world are those with **proprietary insight**: specific, first-hand knowledge of customers that competitors and models do not have. Every real interview you run builds that asset.

Two more shifts to notice:

1. Your customers increasingly ask AI assistants for recommendations, so understanding *how they search and decide* is part of discovery. Add the question: "How did you look for solutions last time?"
2. Customer data you collect is now more valuable and more regulated. Treat it as an asset you protect.

Check your understanding

1. Rewrite "Would you use an app that reminds you to reorder supplies?" for a café-owner interview, and say which claim your rewrite can produce evidence for.
2. A seller says: "Honestly, I'd definitely try something like that." Tag it (claim, E-level, source details) and say what it can and cannot show.
3. Twelve minutes into an interview you described your idea. The person then says they have been paying a bookkeeper 200 CU a month for two years to fix exactly this. How do you tag the two parts of the interview?
4. Your AI rehearsal persona told you "the biggest pain is that suppliers never answer the phone". Can you put this in your ledger? What can you do with it?
5. An interviewee answers "I usually check the orders every evening." What is your next question, and why?

Model answers are in the facilitator guide and open in the workspace after your own attempt.

How this lesson is reviewed

The sprint rubric scores your interview guide and your claim-specific tagging; those scores feed the research and evidence dimensions of the Stage 2 review, and your AI use note feeds the AI accountability dimension. Each founder conducting interviews personally is part of the founder-contribution dimension. Your interview count goes only into your learning result; it is never a rubric criterion.

Sources: Fitzpatrick, *The Mom Test* (2013), principles taught in Foundiro wording, with no rule wording taken from the book; Aulet, MIT OCW 15.390 (2013), lecture 6; Stanford d.school, *Design Thinking Bootleg* (2018), interview preparation and interview-for-empathy cards; Kaufman, "Feedback", personalmba.com; Blank, Customer Development Manifesto (2012); Osterwalder et al., *Value Proposition Design* (2014), The Journalist; the Foundiro Customer Evidence Sprint and Method Rulebook.