How the sample lesson fits the method

This folder holds one complete lesson from Foundiro, exactly as it is taught: **Stage 2, Lesson U2.3, Interviewing without bias.** It is the lesson where most founders first meet the difference between being encouraged and learning something true. This note explains the method behind it in two pages, so a reviewer can see what each piece of the lesson is for.

File	Who it is for	What it holds
A. Method note (this file)	Reviewers and staff	The rules the lesson applies
B. Participant lesson	Students	The full lesson: why it matters, the lesson, a worked example, practice, the application to their own venture, AI support and knowledge checks
C. Participant worksheets	Students	The interview guide, a filled example, the evidence record, the consent record and the AI use note, with writing space
D. Facilitator guide	Teaching staff	Session plan, failure modes, model answers, the rubric, oral questions and a worked review decision
E. Teaching slides	Teaching staff	The slides used in the session

Where the lesson sits

Stage 2 asks one question: does a defined group of people really have this need, and what do they do and spend about it today? The stage runs over about four weeks and ends with a human review (G2, Customer value investigated). Lesson U2.3 is the instrument for the whole stage: if the interviews are biased, every later decision is built on politeness. It is taught inside a connected sequence called the Customer Evidence Sprint, alongside recruiting people ethically (U2.4) and asking about money (U2.7).

The rules this lesson applies

1. Write the claim and the criterion first. Before any interview, the founder writes the claims under test (C1 the need exists for a named group; C2 it matters enough to act on; C3 people already spend money, time or effort on a workaround) and, for each, what result would count as support and what would count against. A criterion written after seeing the answers is allowed but marked "after the data", and it lowers confidence.

2. Ask about the past, not the future. Questions about what people would do produce answers that feel like validation and are worth very little. Questions about the last specific occasion produce stories about what people actually did. The lesson teaches five interview rules and a bad-to-good rewrite routine.

3. Record every item against a claim, with its commitment level and source. Each meaningful statement or observation goes into the evidence record within an hour: which claim it bears on,

how much commitment it shows on the E-ladder (E5 money or repeat behaviour down to E1 opinion, and E0 for anything synthetic), and whether it was observed, reported, a document or a third-party source. The level describes the item. It never decides on its own whether a claim is supported; that judgement weighs relevance, source and the pattern across items.

4. A slip is recorded, not hidden. Anything a person says after the founder has described their idea counts as stated intent at most, unless it comes with a real commitment. The founder notes the minute it happened.

5. Consent is checked first, and it has five parts. Each interview records five separate permissions, each yes, no or pending with scope and date: research notes, recording, AI transcription, external AI analysis and public quotation. An item without consent is excluded. Consent is an integrity condition checked before anything else at the review; it is never scored.

6. AI helps, under rules. The founder drafts the interview guide first. Your AI Partner may then flag leading, hypothetical and pitching questions and play a practice interviewee. The founder keeps or rejects each suggestion and writes a short AI use note. A practice persona is synthetic (E0) and is never logged as evidence. Every step has a route without AI.

What the reviewer checks at G2

The mentor reviews the fixed submission in four steps: integrity (consent, permissions, honest sources), the stage test (are C1 and C2 supported with at least moderate confidence for a named group, against the criteria set in advance), the venture status (GO, CONDITIONAL, NOT YET, or NOT SUPPORTED: PIVOT OR STOP, with what was checked), and the learning result for each founder, recorded separately. The number of interviews completed is a learning requirement; it can never push a venture status up or down.

What makes this lesson different

The lesson is not a talk about customer discovery. Students leave the session with a guide they wrote, tested against an AI critique they had to accept or reject, and they run real interviews the same week. Their evidence is recorded in a form a mentor can check line by line, and the AI's contribution is visible. The worked review in the facilitator guide shows a venture whose first claim set was not supported, and why that was a good result.

Sources credited in the lesson: Rob Fitzpatrick, The Mom Test (principles taught in Foundiro wording); Bill Aulet, MIT 15.390; Stanford d.school, Design Thinking Bootleg (paraphrased); Josh Kaufman, The Personal MBA; Steve Blank, customer development; Osterwalder et al., Value Proposition Design. All cases are fictional.