
Module 3: Problem Discovery and Market Research

Find problems worth solving: research methods, interviews and surveys, personas, jobs-to-be-done, journeys, synthesis and evidence quality.

Product Management Master Handbook
v0.1 · 23 October 2026

Built from the 55-lecture course study guide and extended to the 12-module Product Core curriculum.
Hypothetical examples are labelled; no real company results are invented.

Contents

3.1 What discovery is	3
3.2 Research methods: qualitative and quantitative, primary and secondary	4
3.3 Customer interviews	5
3.4 Surveys and questionnaire bias	7
3.5 Customer segmentation and personas	8
3.6 Jobs to Be Done (JTBD)	9
3.7 Pain-point identification	10
3.8 Customer journey mapping and user flows	10
3.9 Competitor research, review mining and support-ticket analysis	12
3.10 Behavioural analytics and observation	13
3.11 Research synthesis	13
3.12 Problem validation and evidence quality	15
3.13 Guesstimates (estimation)	15
3.14 Product critique and teardown	16
3.15 Discovery exercises	17
3.16 Source-flagged corrections	18
3.17 Interview perspective	18
3.18 Key takeaways	18
3.19 Check your understanding	20
3.20 Recommended resources	21

MODULE 3

Problem Discovery and Market Research

Find problems worth solving: research methods, interviews and surveys, personas, jobs-to-be-done, journeys, synthesis and evidence quality.

≈ 10 h study

About examples

Personas, quotes, review counts and estimates are **invented** to teach technique. They are not research findings about any real product or user group.

3.1 What discovery is


Product discovery is the work of learning *what is worth building* before (and while) building it: who has the problem, how painful it is, how they cope today, and whether your idea would be chosen and used. **Customer discovery** focuses on understanding customers; **problem validation** tests that the problem is real, frequent and painful enough to matter.

Level 1 · Beginner

Discovery = talking to people and watching what they do, so you build the right thing. If you can only ask "Would you use my idea?", you will mostly hear polite yeses.

From your course notes (Lectures 6, 8; Picture lesson)

Validate the problem before investing heavily in a solution. Ask whether the need is real, who experiences it, how people solve it today and whether the opportunity is worth pursuing. *A solution is a hypothesis to test, not a guaranteed success.* Start with a user and a task; ask about a recent real experience; watch the task where appropriate; look for repeated difficulties. **Do not assume an idea is useful simply because someone says it sounds nice.**



A discovery cycle. Discovery repeats; each loop is cheap relative to building. Most teams under-invest in Synthesise and Decide.

3.2 Research methods: qualitative and quantitative, primary and secondary

Two different axes — *do not confuse them*:

	Qualitative (why / how)	Quantitative (how many / how much)
Primary (you collect it)	Interviews, observation, diary studies, usability tests, moderated focus groups	Surveys, A/B tests, instrumented behaviour in your own product
Secondary (already exists)	Support tickets, app-store reviews, forums, published case studies	Market reports, public statistics, company databases, benchmarks

From your course notes (Lecture 6)

Qualitative and quantitative describe the kind of evidence, not alternatives to primary research. Interviews give qualitative insight into *why*; a well-designed survey can estimate *how common* a behaviour is. Both can be primary. A **focus group** is a moderated research discussion — *not another name for a target audience*. Ten interviews reveal useful clues but do not establish the habits of an entire city.

Question you have	Best-fit method	What you get	What you don't get
Why do users abandon checkout?	Interviews + session observation	Reasons, context, language	How common each reason is
How many abandon at the fee step?	Funnel analytics	Counts and rates	Why
Which of three designs is clearer?	Usability test (5–8 per round)	Where people struggle	Statistical proof
Which price converts better?	A/B or priced pilot	Behaviour at scale	Reasons behind it
How big is the segment?	Bottom-up sizing + survey	An estimate with assumptions	Proof of demand
What do users say publicly about rivals?	Review mining	Themes and vocabulary	Representative opinion

Choosing a method by question. Mix methods: use analytics to find **where**, interviews to learn **why**, experiments to test **whether the fix works** (triangulation).

3.3 Customer interviews

3.3.1 Principles

1. **Ask about the past, not the future.** “Tell me about the last time you ordered lunch” beats “Would you use an app that...?”
2. **Seek stories, then specifics:** what triggered it, what alternatives they considered, what happened, how they felt, what they did next.
3. **Stay neutral:** don't pitch; don't lead; tolerate silence.
4. **Recruit for contrast:** include people who tried and churned, people who never tried, power users — not only fans.
5. **Consent and privacy:** explain purpose, record only with permission, remove identifying details in notes.
6. **Debrief fast:** write observations separately from interpretations within 24 hours.

From your course notes (Lecture 3; Companion ch. 2)

Interview a student who buys lunch. Ask *"Tell me about the last time you ordered lunch."* If they say they missed delivery fees, watch where they expected to see the price. **Five interviews can reveal problems but cannot establish their population frequency.** Recruit a range of users including those who abandoned and those who did not; get consent; avoid leading questions; remove identifying information; compare themes with quantitative data.

3.3.2 Question design

Weak (leading / hypothetical)	Stronger (past behaviour, neutral)
"Would you pay for faster delivery?"	"Last time you were in a hurry, what did you do about food?"
"Do you find our checkout confusing?"	"Walk me through placing your last order. What did you look at first?"
"Don't you hate hidden fees?"	"How did you decide whether the total was acceptable?"
"Is price important to you?"	"Tell me about a time you decided <i>not</i> to order because of the cost."
"What features do you want?"	"What did you do the last time this didn't work?"

3.3.3 Sample interview script (45 minutes)

Script: ordering food after work (hypothetical study)

Opening (5 min). "Thanks for your time. I'm learning how people decide about dinner on weekdays. There are no right answers; I'm not testing you. May I record audio for my notes only? I'll remove names." **Warm-up (5 min).** "Tell me about your typical weekday evening." **Core story (20 min).** "Think about the last time you ordered food from an app on a weekday. What was happening before you opened it?" → "What did you do first?" → "What else did you consider?" → "What were you looking at when you decided?" → "Was there a moment you almost gave up?" → "What did you do next?" **Probe the friction (10 min).** "You mentioned the total changed at the end — what did you expect to see?" → "How did that feel?" → "What did you do?" **Wrap-up (5 min).** "Is there anything I should have asked?" → "Who else should I speak to?" → thank, explain next steps. **Don't:** pitch your idea, ask about features, or suggest answers.

3.3.4 Recording the evidence

Use a two-column note sheet: **Observation** (what they did/said, verbatim where possible) | **Interpretation** (your hypothesis). Tag each observation with the participant ID and the journey stage.

3.4 Surveys and questionnaire bias

A **survey** estimates *how common* something is — if the sample is representative and the questions are neutral.

Common biases and fixes

Bias	What happens	Fix
Leading question	Wording nudges the answer ("How great was...?")	Neutral wording; balanced scales
Double-barrelled	Two questions in one ("fast and cheap?")	One idea per question
Social desirability	People give flattering answers	Anonymity; behaviour ("last week, how many times...?")
Acquiescence	Tendency to agree	Mix positively and negatively worded items; use specific choices
Order effects	Earlier questions prime later ones	Randomise; put sensitive items last
Sampling / self-selection	Respondents differ from the population	Define the frame; track response rate; compare with known figures
Hypothetical bias	Stated intent $\neq$ behaviour	Prefer behavioural questions; validate with a real action
Incentive bias	Rewards attract the wrong people	Screen carefully; use modest incentives

From your course notes (Lecture 6; Lecture 10)

*The claim that "90% of survey responses are biased" was not substantiated; sampling, incentives and question design can introduce bias but do not make surveys inherently invalid. **A survey preference is not proof that respondents will purchase.***

Sample survey questions (rewritten)

- Bad: "Do you agree our app is easy and fast to use?" → Good: "In the past 7 days, how many times did you place an order? (0 / 1 / 2–3 / 4+)" and "How easy or difficult was it to find the total price on your last order? (Very difficult ... Very easy)".
- Bad: "Would you use a feature that tells you about fees?" → Good: "When you last saw a price at checkout, was the total different from what you expected? (Yes, higher / Yes, lower / No / Don't remember)".

3.5 Customer segmentation and personas

A **segment** is a group sharing a *relevant* need or behaviour. A **persona** is an illustrative profile of a segment — a communication aid, not proof everyone acts alike. Prefer **behavioural and need-based** segments over purely demographic ones.

From your course notes (Lecture 3)

A target audience is a group a product or message is intended to reach. **The buyer and end user can be different** (FirstCry: a parent may buy while a child uses). Build and validate personas with research, *not guesses*. **There is no universal rule that every product needs two or three personas**, and choosing the largest segment alone does not prove product–market fit. The iPhone-user personas in class were exercise assumptions, not claims about all iPhone users. When needs conflict, consider market size, severity of need, strategic fit, feasibility and reachability; persona overlap is normal.

NAME & CONTEXT

- "Asha", 31, product analyst
- Works late 3–4 days a week
- Uses food apps on her commute

GOALS & MOTIVATIONS

- A reliable dinner without planning
- Keep weekday spend predictable

BEHAVIOURS (FROM INTERVIEWS)

- Re-orders favourites; checks total before paying
- Abandons if total jumps at last step

PAINS / FRUSTRATIONS

- Fees appear late
- Delivery estimates drift
- Unclear refund path

NEEDS FROM THE PRODUCT

- Clear total early
- Trustworthy ETA
- One-tap reorder

EVIDENCE & CONFIDENCE

- 8 interviews (Mar 2026, hypothetical)
- Confidence: medium
- Open questions: what happens at month-end?

Persona template filled with a hypothetical example. A persona is only as good as its evidence. Always record sample size, date and confidence, and update it when research changes.

3.6 Jobs to Be Done (JTBD)

JTBD frames demand as **people "hiring" a product to make progress in a particular circumstance**. A **job statement** has the form:

When [situation], I want to [motivation / progress], so I can [desired outcome].

- Jobs have **functional**, **emotional** and **social** dimensions.
- The *competition* is whatever else gets hired for the job — including doing nothing.
- Uncover jobs from **switching stories**: what triggered the move, alternatives considered, what pushed them away from the old solution (push), what pulled them to the new one (pull), what anxieties held them back, what habits kept them in place — the **four forces**.

Example (hypothetical): lunch

When I'm in back-to-back meetings and have 15 minutes, I want to get a filling lunch without leaving my desk or being surprised by the cost, so I can stay focused and not blow my budget.

Competing alternatives: order in, ask a colleague to pick up, eat a snack, skip lunch, bring food. **Anxieties:** "Will it arrive in time? Will the price change?" **Implication:** reliability and total price clarity matter more than restaurant variety for this job.

PUSH (AWAY FROM THE CURRENT WAY)

- Wasted groceries
- Planning takes an hour

PULL (TOWARD THE NEW SOLUTION)

- Weekly plan with a price estimate

ANXIETY (ABOUT THE NEW SOLUTION)

- Will estimates be accurate? Will it fit our diet?

HABIT (OF THE PRESENT WAY)

- Family's recipe notebook; saved lists in notes app

The four forces of switching (hypothetical switch to a new meal-planning app). A product wins when push + pull outweigh anxiety + habit. Discovery asks about each force; design and onboarding reduce anxiety.

From your course notes (Lecture 21)

Jobs to be done asks what *progress users want to make in a situation*; first-principles thinking breaks a problem down to fundamental assumptions. The recording mainly refers learners to further material for these concepts — this section is independent reference material.


3.7 Pain-point identification

List pains with **severity** (how bad), **frequency** (how often), **breadth** (how many people) and **current workaround** (the clearest sign of real pain: people already spend effort to solve it). Rank by severity × frequency × breadth, then by whether you can credibly solve it.


3.8 Customer journey mapping and user flows

From your course notes (Lecture 5; Companion ch. 3)

A **journey map** follows the *whole experience*, including actions, thoughts and feelings, across touchpoints; a **user flow** focuses on paths and decisions inside an interface; one journey can contain many flows. A **workflow** focuses on the sequence of activities, systems and decisions. Start with a concrete task, map the steps and important branches; use the appropriate level of detail. AI can brainstorm a first draft (give it problem, users, features, constraints) — **check it against the real product and do not copy it blindly**. Add the key branches: valid/invalid address, payment failure, retry, cancellation — and label the desired outcome for each. A remembered login session is not necessarily single sign-on.



Journey: weekday dinner order (hypothetical, one persona). Journey stages are the rows of a map; add rows for actions, thoughts, emotions, touchpoints, pain points and opportunities.



Emotion curve across the journey (hypothetical, from 8 interviews). Hypothetical. The two lowest points — the cart total and payment — point to where discovery and design effort should focus.

3.8.1 Opportunity Solution Trees (OST)

Teresa Torres's **Opportunity Solution Tree** connects a **desired outcome** to the customer **opportunities** (needs, pains, desires) and then to **solutions** and **experiments**. It prevents jumping from outcome straight to a favourite feature.



Opportunity Solution Tree (hypothetical). Each solution should be paired with the cheapest experiment that tests it. Compare opportunities (not just solutions) when you prioritise.


3.9 Competitor research, review mining and support-ticket analysis

Review mining turns public reviews and support tickets into themes and a vocabulary.

Method: (1) sample reviews across ratings and dates; (2) code each into themes (one review can have several); (3) count and read representative quotes; (4) note sampling limits; (5) turn themes into hypotheses.

Worked example — coding 40 synthetic app reviews

Theme	Mentions	Share of 40 reviews
Late delivery	11	27.5%
Fees / total unclear	9	22.5%
Food quality	7	17.5%
App crashes at payment	6	15.0%
Slow support	4	10.0%
Good discounts (positive)	3	7.5%



Themes in 40 synthetic reviews (hypothetical). Synthetic data for teaching coding. With 40 reviews you can form hypotheses, not estimate the share of all customers. Reviews over-represent extreme experiences.

From your course notes (Lecture 39)

The class sampled a public dataset, asked an AI tool for sentiment and potential themes (reliability, pricing, safety, support). **Review individual comments and sampling bias; a small sample cannot establish the full market's sentiment.** Public datasets can be real, synthetic or mixed — check provenance and licence. Never upload confidential data without authorisation.

3.10 Behavioural analytics and observation

- **Behavioural analytics** (funnels, cohorts, session replays where permitted) shows *what* people do at scale (Module 7).
- **Observational research** — watching users perform tasks (in context if possible) reveals what they cannot articulate: workarounds, hesitations, misreadings. Think-aloud usability testing is a form of observation.
- Always combine: analytics finds *where*; observation and interviews find *why*.

3.11 Research synthesis

3.11.1 Affinity mapping

Write each observation on a note; cluster by similarity (silently first, then discuss); name clusters as **insights** (“Users distrust totals that change at the end”), not topics (“Pricing”); count participants per cluster.

INSIGHT A · 6 OF 8

- "The total jumped at the end."
- "I didn't see delivery fee until payment."
- "I wasn't sure what I was paying."

INSIGHT B · 4 OF 8

- "If payment fails I start over."
- "I switched to cash once."


INSIGHT C · 3 OF 8

- "I just want to reorder my usual."
- "Saved addresses save me time."

Affinity map output (hypothetical, 8 participants). Hypothetical. Report counts as "6 of 8 participants", never percentages from a handful of interviews.

3.11.2 Assumption mapping and hypothesis generation

List every assumption behind the idea (desirability, usability, feasibility, viability). Plot on **importance** (how fatal if wrong) vs **evidence** (how much we know). Test the **important, low-evidence** assumptions first.



Evidence we have (1 = none, 10 = strong)

Assumption map (hypothetical meal-planning app). Hypothetical. Top-left (important, little evidence) is where experiments should start.

Hypothesis template: We believe that [change/solution] for [segment] will [expected outcome/metric] because [reason/evidence]. We will know we are right when [measurable signal].

3.12 Problem validation and evidence quality

Triangulation means reaching a conclusion from multiple independent kinds of evidence.

Grade evidence honestly:

Strength	Evidence	Example
Strong	Observed behaviour at scale; randomised experiments; paid pilots	Conversion lift in an A/B test; people paying a real price
Medium	Repeated behaviour in small samples; consistent interview themes; usability test patterns	6 of 8 users miss the fee in a task
Weak	Stated intent, opinions, single anecdotes, AI-generated claims, "I would use this"	Survey says 80% "would try"

From your course notes (Lectures 6, 28, 41)

Label comparisons based on impressions as **hypotheses**; support material claims with suitable research, sources and methodology. **Do not present AI-generated target numbers as evidence** — one team acknowledged the tool invented its metrics. **A student's own experience is a starting point, not proof of a widespread problem.**

3.13 Guesstimates (estimation)

A **guesstimate** breaks an unknown number into smaller numbers you can explain, state assumptions for, and sanity-check — and gives a **range**.

From your course notes (Lecture 7)

Clarify the question first (geography, typical day vs special event, ride types, whether cancelled requests count — *keep definitions consistent*). Break into measurable assumptions; **explain and sanity-check** each; compare **top-down and bottom-up**; communicate a range and the assumptions that matter most. The class estimates were **practice figures, not verified Uber statistics**. *There is no single right answer, but arithmetic errors, inconsistent definitions or implausible assumptions make an estimate weak.* Simple example: 1,000 students × 60% at school × 50% buying lunch = 300 lunches.

Worked example — daily ride-hailing trips in a large metro (practice numbers, not facts)

Top-down: metro population 20 million → adults 70% = 14.0 million → people who use ride-hailing at least weekly: 8% = 1.12 million → 3 trips per week each = 3.36 million trips/week → **≈ 480,000 trips/day**. **Bottom-up (supply check):** active drivers per day 55,000 × 9 completed trips per driver per day = **≈ 495,000 trips/day**. **Reconcile:** two independent approaches land within ~3% — *this raises confidence but both rest on assumptions I invented*. Give a range: **450k–500k/day**, and name the most uncertain assumptions (share of weekly users, trips per driver). Sanity check: 480k trips/day ÷ 14M adults ≈ 3.4% of adults ride each day — plausible for a big city.

Structure to use in interviews: clarify → choose approach (top-down or bottom-up) → segment → state assumptions → compute → sanity-check → give range → discuss what would change it.

3.14 Product critique and teardown

A **teardown** studies a product from the outside to understand users, goal, journey, strengths and possible improvements. Separate **observations** from **assumptions**; an outsider cannot know a company's internal data.

From your course notes (Lectures 4, 19, 28)

A critique explains **the user, the task, the difficulty and the expected benefit**. Suggested teardown sequence: product context and problem → vision/purpose → personas → user and market research → journey and pain points → prioritised opportunities → proposed solution → success metrics (adapt it to the question). Each claim needs evidence and each change needs a test. Student proposals in class (Myntra, Amazon, X, ChatGPT, Netflix) were *proposals and observations, not verified gaps*. Consider constraints such as content rights and contracts.

Teardown skeleton (use for any product)

1. **Product & user:** what job; for whom; business model (assumption, label it).
2. **Journey:** 5–7 stages; where does it help; where does it hurt (evidence: your own walk-through, reviews, tests).
3. **Opportunity:** one problem; evidence; who is affected; size (assumption).
4. **Solutions:** 2–3 options incl. a non-AI/simple one.
5. **Choice:** impact/effort/risk; recommendation.
6. **Measurement:** primary metric; guardrails; experiment.
7. **Risks and unknowns.**

3.15 Discovery exercises

Exercise 1 · Rewrite the questions

Rewrite three poor interview questions and explain the bias in each: (a) "Do you think our onboarding is too long?" (b) "Would you pay ₹99/month for this?" (c) "What features would make you love this app?"

Model answer

(a) Leading and closed → "Walk me through your first time setting up the app. Where did you hesitate?" (b) Hypothetical bias → "What do you pay for today to solve this? How did you decide it was worth it?" (c) Solution-seeking; customers aren't product designers → "Tell me about the last time this task went badly. What did you do?"

Exercise 2 · Evidence grading

You have: (1) 8 interviews, 6 mention surprise at fees; (2) a funnel showing a 30% drop at the fee step; (3) a survey where 78% say they "like transparent pricing". Rank the evidence and state what you would do next.

Model answer

(2) funnel = behavioural, scalable (strong *that* a drop exists, not why). (1) interviews = medium (consistent reasons, small sample). (3) survey = weak (stated preference, hypothetical). Next: segment the funnel (device, new/returning), run a usability task on the fee screen, then test early fee display in an A/B experiment with a complaint guardrail.

Exercise 3 · Estimate

Estimate how many cups of coffee are sold in a city of 5 million in a day. Show two methods and a range.

Model answer (example)

Top-down: $5M \times 70\% \text{ adults} = 3.5M$; $55\% \text{ drink coffee} \approx 1.9M$; $1.8 \text{ cups/day from cafés/vending/home?}$ Only 35% bought outside $\rightarrow 1.9M \times 1.8 \times 0.35 \approx \mathbf{1.2M \text{ purchased cups/day}}$. **Bottom-up:** $2,500 \text{ cafés} \times 250 \text{ cups} = 625k + 1,500 \text{ kiosks} \times 300 = 450k + \text{vending/office } 100k \approx \mathbf{1.18M}$. Range 0.9–1.4M. Sensitive to share who buy outside and cups per café.

3.16 Source-flagged corrections

- Qualitative/quantitative $\neq$ primary/secondary.
- Focus group $\neq$ target audience.
- Persona count is not fixed; personas must be evidence-based.
- “Surveys don't work in some country” is unsupported; design matters.
- All summaries of research need source checking.

3.17 Interview perspective

“How would you validate a problem before building?” — 60-second answer

“I'd start with a clear problem hypothesis and target segment, then interview 8–10 people who've recently experienced it, asking about their last real experience and what they do today. I'd look at available data to size and locate the problem, mine reviews or support tickets for patterns, and write down the assumptions that could kill the idea. For the riskiest assumption I'd run the cheapest test — a prototype or a smoke test with a real price — and decide in advance what result would make me continue, pivot or stop.”

“Tell me about a product you love and how you'd improve it.”

Tests: product sense, empathy, prioritisation, honesty about evidence. *Structure:* user + job $\rightarrow$ why it works $\rightarrow$ one meaningful pain (with evidence you'd gather) $\rightarrow$ 2–3 solutions $\rightarrow$ pick one with trade-offs $\rightarrow$ metric + guardrail + experiment. *Mistakes:* naming five features, ignoring constraints (content rights, privacy, cost), assuming your own experience is universal.

3.18 Key takeaways

- Discovery reduces **risk** by learning *before* building. Treat solutions as hypotheses.

- Ask about **past behaviour**; recruit for contrast; separate observation from interpretation.
- Match method to question; **triangulate** (analytics → interviews → experiment).
- Surveys estimate frequency but suffer from bias — design carefully.
- Use **JTBD**, journey maps, OSTs and assumption maps to structure what you learn.
- Grade evidence honestly; label hypotheses and AI-generated claims.

3.19 Check your understanding

CHECK YOUR UNDERSTANDING

Which is an example of secondary quantitative research?

- A. A survey you design and send
- B. A published market report with spending figures
- C. Five customer interviews
- D. A usability test

Why: Secondary research uses existing information; quantitative means numeric evidence. A published report fits both.

Which interview question is best?

- A. Would you use an app that shows fees early?
- B. Don't you hate surprise fees?
- C. Tell me about the last time the total changed at checkout.
- D. Which features do you want?

Why: It asks about a specific past experience, is neutral and invites a story.

In the four forces of switching, 'anxiety' refers to...

- A. Frustration with the current solution
- B. Attraction to the new solution
- C. Worries about trying the new solution
- D. Inertia of habit

Why: Push and pull move people toward switching; anxiety and habit hold them back.

6 of 8 interviewees mention the fee surprise. How should you report it?

- A. 75% of customers dislike hidden fees
- B. 6 of 8 interviewed participants mentioned it; frequency in the population is unknown
- C. It's statistically significant
- D. Ignore it; the sample is too small

Why: Small qualitative samples reveal themes, not population percentages.

On an assumption map, which should you test first?

- A. High importance, low evidence
- B. Low importance, high evidence
- C. Low importance, low evidence
- D. High importance, high evidence

Why: Assumptions that would sink the idea and that you know little about carry the most risk per test.

A guesstimate gives 480k/day top-down and 495k/day bottom-up. The right conclusion is...

- A. The answer is exactly 487,500
- B. Confidence is higher but both depend on assumptions; give a range
- C. Both are wrong
- D. Use the larger number

Why: Agreement of independent methods raises confidence, but communicate a range and the key assumptions.

3.20 Recommended resources

- [Nielsen Norman Group: Heuristic evaluations](#) (from your course list) and [user interviews](#).
- [Product Talk: Opportunity Solution Trees](#).
- Rob Fitzpatrick, *The Mom Test* — asking good questions about past behaviour ([book site](#)).
- [Jobs to Be Done — Christensen Institute](#).
- [Google Cloud / Google: Gen AI glossary](#) for AI-assisted research caveats.