AI Personal CFO
A design for every screen

Twenty-five screens, redrawn in the design language this repository already has, with the reason for every decision.

Evidence: `docs/screens/SCREENS.pdf` — the 115 committed renders of what ships today.
Language: `docs/design/design-language-v2.html` — the v2 spec and its seven moves.
Mockups: `docs/design/redesign/mockups/` · rendered to `shots/`.
No application code was changed to produce this. Nothing in `:feature:*`, `:core:*` or the build was touched; this is a proposal to argue with, not a patch.

How this was made

Three inputs, in this order.

- **What ships today.** The 115 Paparazzi baselines, read side by side in `SCREENS.pdf`. Every criticism below points at a specific render, not at a memory of one.
- **The language the repository already chose.** `docs/design/design-language-v2.html` defines v2 — the trust-teal seed, the token table, and seven moves — on six screens. **This pass adds no eighth move and no new token.** It extends those six to twenty-five, which is the work v2 left undone.
- **The rules that outrank taste.** `.claude/skills/screen-design/SKILL.md` lists the non-negotiables: colour never signals alone, no figure without provenance within one tap, every string a resource, money through `MoneyFormatter`, no fixed heights, demo data only. A design that breaks one of those is wrong here however good it looks.

Each screen below gets the same treatment: the shipped render, the proposed one, what changed, and why — with the move, the standard and the project rule that forced it named, so a reviewer can disagree with a reason rather than with a preference.

> **What is out of scope, deliberately.** No new engine, no new dependency, no new colour seed, no new font, no change to what any figure *is* — every number in these mockups is the same engine result the app already computes (P-03). The redesign moves hierarchy, order and evidence; it does not move arithmetic.

Coverage — all 25 shipped screens

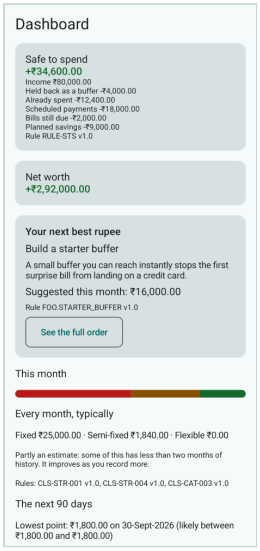
The app ships 25 screens, collected in `docs/screens/SCREENS.pdf`. This document designs every one of them: no screen was dropped as too small, too rare or too awkward, because the screens that get skipped in a redesign are exactly the ones that stay ugly.

Shipped screen (v1)	v2 mockup	Argued in
1. Onboarding	`onboarding`	§1
2. App lock	`lock`	§2
3. Dashboard	`today`	§3
4. Order of operations	`order`	§4
5. Add transaction	`add`	§5
6. Transactions	`txns`	§6
7. Receipt review	`receipt`	§7
8. Bank-alert drafts	`sms`	§8
9. Categories	`categories`	§9
10. Accounts	`money`	§10
11. Account editor	`accountEditor`	§11
12. Holdings	`holdings`	§12
13. Allocation	`allocation`	§13
14. Budgets	`budgets`	§14
15. Goals	`goals`	§15
16. Goal detail	`goalDetail`	§16
17. Emergency fund	`emergency`	§17
18. Purchase advisor	`advisor`	§18
19. Simulators	`simulators`	§19
20. Opportunity	`opportunity`	§20
21. Vehicles	`vehicles`	§21
22. Assistant	`ask`	§22
23. Settings	`settings`	§23
24. Consents	`consents`	§24
25. Erase everything	`erase`	§25
25 screens	**25 designs**	**25 sections**

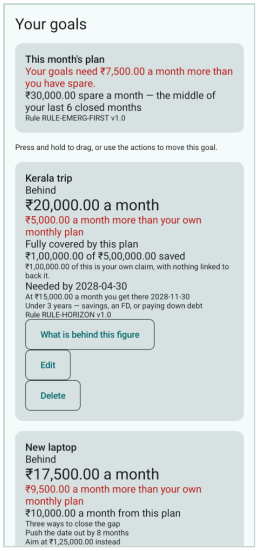
> **There is a second proposal.** `docs/design/AI_Personal_CFO_v3.pdf` answers the same 25 screens from a different premise — the app as one queue of decisions rather than a set of dashboards — in this same palette. The two are meant to be read against each other; v3's closing section says how to choose between them.

The diagnosis

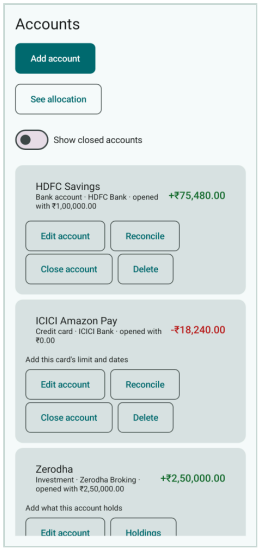
Seven problems, each visible in more than one shipped screen. They are symptoms of one cause: the screens were built as faithful renderings of their `UiState`, so everything the state knows is on the screen at the same weight.



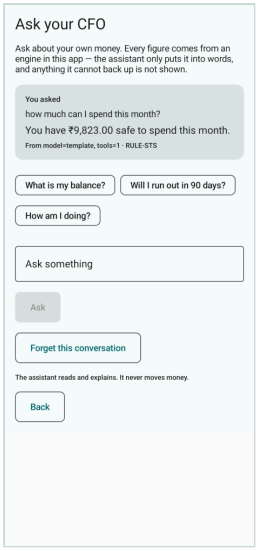
Dashboard: six ledger lines under the one figure that matters



Goals: ten lines of prose, three ragged buttons, no bar



Accounts: four full-width buttons per row



Ask: a conversation with no conversation layout

1 · No screen says which figure it exists for

The dashboard's Safe-to-Spend sits in the same card as the six amounts it was derived from, in roughly the same size. A goal's name, its status, its required monthly and its rule citation are four consecutive lines at four similar weights. When everything is emphasised nothing is, and the user is left to do the ranking the app was supposed to do for them.

2 · Provenance is printed as body text

`Rule RULE-STS v1.0`, `Rule FOO.STARTER_BUFFER v1.0`, `From AI-VEH v1.0 · VEH-KB.service` — on nine screens, in the content column, at content size. P-02 is right that the working must be available; printing it inline makes every card 20% longer and teaches users to skip text. **The promise is "show the work", not "show the work at all times".**

3 · Actions are stacked, ragged and undifferentiated

An accounts row carries Edit, Reconcile, Close and Delete as four equal outlined buttons; a goal carries three; the add-transaction screen ends with Save, Cancel, "Scan a receipt" and "From your bank messages" at four different widths. Destructive actions sit at the same weight as routine ones, and Material's own guidance — one emphasised action per view — is nowhere.

4 · Quantities are described where they could be drawn

"₹1,00,000.00 of ₹5,00,000.00 saved", "80% used", "Spent this month ₹4,600.00" — all correct, all text. The dashboard's one bar has no labels, no legend and no scale, so it says less than the sentence beside it.

5 · Colour and raw data do work the words should do

Red is used for "needs more than you have spare" (an attention state, not an error), so the one colour that should mean *stop* appears on a routine planning screen. Dates print as `2028-04-30`; engine ids print verbatim; a rate reads `36.00% a year` beside a figure in paise.

6 · Navigation is a button at the bottom of a scroll

Screens end with *Back*, *Done* or *OK* buttons rather than carrying a top app bar with a back affordance, and the app has no persistent navigation at all — every destination is reached from a list on the dashboard. Android users expect a bar; the platform draws one for free.

7 · Empty, blocked and disabled states explain nothing

"Ask the advisor" and "Ask" are rendered disabled with no statement of what is missing. The SMS screen with no consent, the allocation with nothing priced and the simulators with no debts each say one short sentence where the user needs to know what to do next.

Foundations — what this design inherits

v2 is not a repaint and this pass does not make it one. The Material 3 trust-teal seed, the semantic roles, Roboto and the existing `:core:designsystem` tokens all stay exactly as they are. What follows is the inherited part, restated so the screens can be read against it.

The seven moves, applied in order

Move	What it means when drawing a screen
1 · One answer	Name the single figure the screen exists to deliver and give it the largest type. If you cannot name it, the screen is not designed yet.
2 · A meaning line	Under every figure, a sentence built from the same engine result. `₹18,420.00` is data; "₹1,180 a day until 31 October, after rent and EMIs" is an answer.
3 · The work is one tap	An inline disclosure on the card: inputs, rule id *and version*, engine, and when it ran. Never a separate evidence screen — v1 had one and nobody reached it.
4 · Pace is drawn	One bar carries spent, the even-pace mark and the overshoot. The projection stays in words, because that is the part a person acts on.
5 · One primary action	One filled button per screen; everything else is text. Chips at 44dp, never a dropdown where chips fit.
6 · Trust in the chrome	`offline · all local` in the status line; consent rows showing both dates; erase stating what it cannot reach.
7 · Calm surfaces	Two surface levels, one hairline instead of shadows, 20dp card / 16dp button / 999dp chip, amounts in tabular figures larger than anything near them.

Tokens — unchanged, and that is the point

primary #00696E — one action per screen

canvas #EEF3F2

surface #FFFFFF — cards holding a figure

surface-2 #F6FAF9 — quieter cards

positive #146C2E — always with “+”

negative #BA1A1A — always with “−”

warning #8A5000 — always with a word

Type: figure 34sp · figure-small 19sp · title 20sp · body 15sp · meaning 13sp · label 11.5sp at 0.06em. Spacing 4 / 8 / 12 / 14 / 20. Radii 20 / 16 / 999. Every amount uses Indian grouping and tabular figures through `MoneyFormatter` (MNY-001) — never a hand-rolled format, never a `Double`.

Navigation — five tabs, and links that hang off what they explain

Today, Money, Plan, Ask, You. The shipped app has no persistent navigation: every destination is a row on the dashboard, which is why that screen became a directory. Five tabs is the Material 3 navigation-bar maximum, and the grouping follows what a person is doing rather than how the modules are split:

Tab	Holds	Why these together
Today	Safe-to-Spend, the insight worth acting on, the next best rupee → the full order of operations	The one question asked daily: can I spend, and what should the spare do.
Money	Accounts, one account → holdings → allocation, the ledger, add / receipt / bank-alert drafts, categories	Everything that is a record of what happened.
Plan	Budgets, goals → one goal, emergency fund, the advisor, the simulators, vehicles, opportunity	Everything that is about what happens next.
Ask	The assistant	Its own tab because it answers across all of the above, and because a conversation keeps state.
You	Money plan, app lock, language, backup, consents, erase	Settings and the privacy promises, which are the same subject here.

Two rules from the skill apply and are easy to break by accident: **a link hangs off the thing it explains** — the pace bar opens budgets, an account row opens that account — and **never a shelf of equally weighted chips**, which is the failure v2 exists to fix. Every screen in this document is reachable within three taps of a tab.

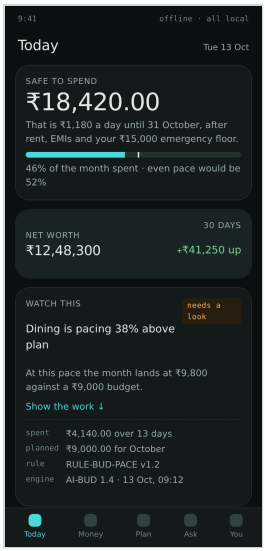
Components added in this pass, and why each was unavoidable

Component	Why no existing v2 component could do it
Sub app bar (back + title)	v2's `.appbar` is a tab root and has no back affordance. "A stack per tab" needs one that does.
Status pill	`.tag` is a provenance mark and reads as metadata. A verdict — *behind*, *over the cap*, *75 days overdue* — needs to read as state. Always a word, never a colour alone.
Segmented control	For exactly-one-of-two-or-three properties of the thing being created (expense/income/transfer, asset/loan/card). M3 uses chips for filters and segments for opposites; the shipped radio buttons are neither.
Stage list	§36's claim is that the *order* is the advice. A row list loses the sequence; the stage list keeps it with a status dot per step.
Leading glyph	A ledger is scanned by shape before it is read. Two letters and a tone, not an icon font — a mock must not promise an asset the app does not bundle.
Banner	For a condition the screen is *in* (a duplicate receipt, nothing consented, this cannot be undone). A card would make it look like content.
Donut + legend	The allocation screen's one answer is a proportion, and a proportion is the one thing a bar cannot say at a glance. Legend text carries the numbers, so the chart is never the only source.

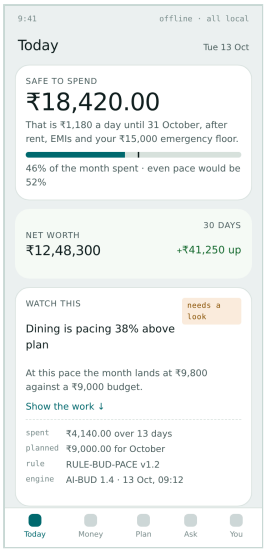
The non-negotiables, as they appear in these screens

- **Colour never signals alone.** Every pill carries a word; every amount carries `+` or `−`. The designs survive greyscale, colour-blindness and the privacy blur — and the blur is why a figure's meaning can never live in its colour.
- **Provenance within one tap.** Every card holding a figure has *Show the work*, and the disclosure names the rule id, its version, the engine and when it ran (AI-ARC-003). It is one tap, not zero, and not two.
- **No fixed heights.** Nothing here is sized to its content's current length, so the designs hold at 200% font and in dark mode — both of which are screenshot-tested, so a design that cannot hold is a design that fails the build.

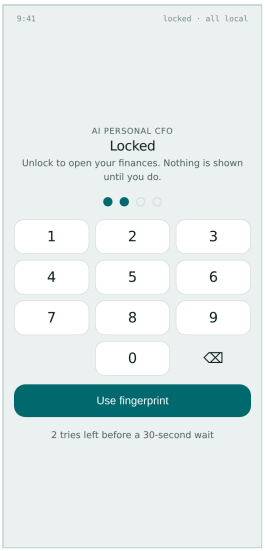
- **Demo data everywhere.** Every figure in these mockups is seeded, and consistent with the v2 prototype's one ledger, so no screen contradicts another.



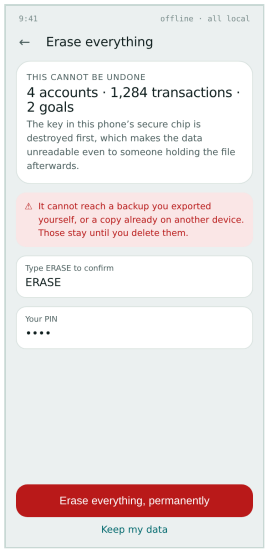
The same screen, dark. Tokens flip; the hierarchy does not.



Today, light



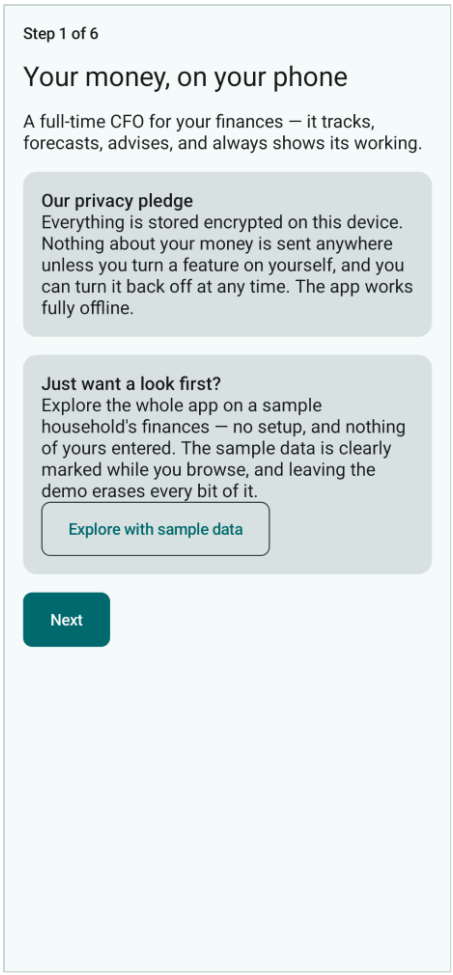
The lock: a keypad, not a text field



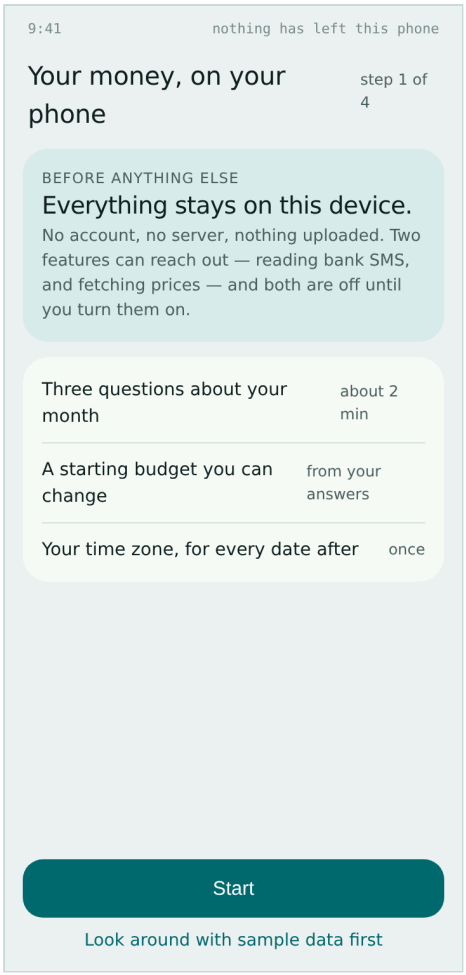
Erase: the one destructive primary in the app

1 · Onboarding

:feature:onboarding · FR-ONB-001/002/003 · moves 1, 5, 6



Today — the shipped render



Proposed — v2

What changed

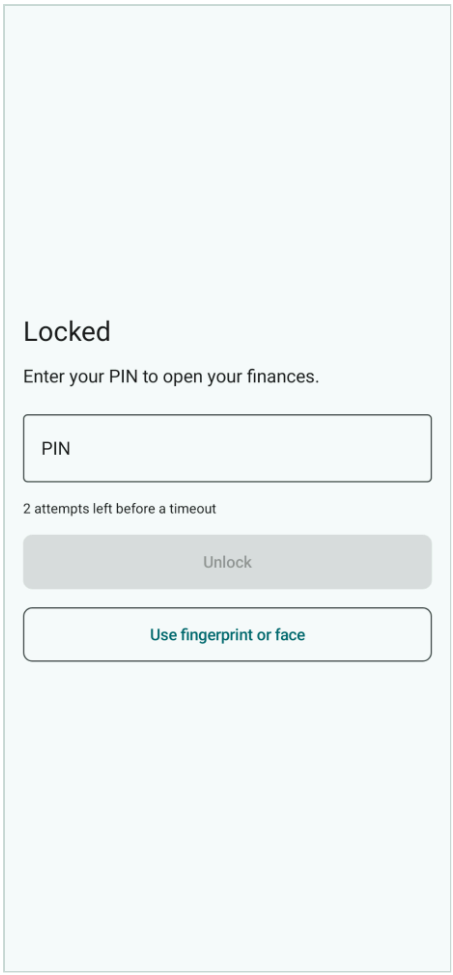
- The privacy pledge becomes the screen's one answer — *Everything stays on this device* — and the two features that *can* reach out are named in the same breath, both off.
- The sample-data offer drops from a second card to a text alternative under one filled button.
- The step counter moves into the app bar; the body lists what the four steps will ask for and what it costs in time.

Why

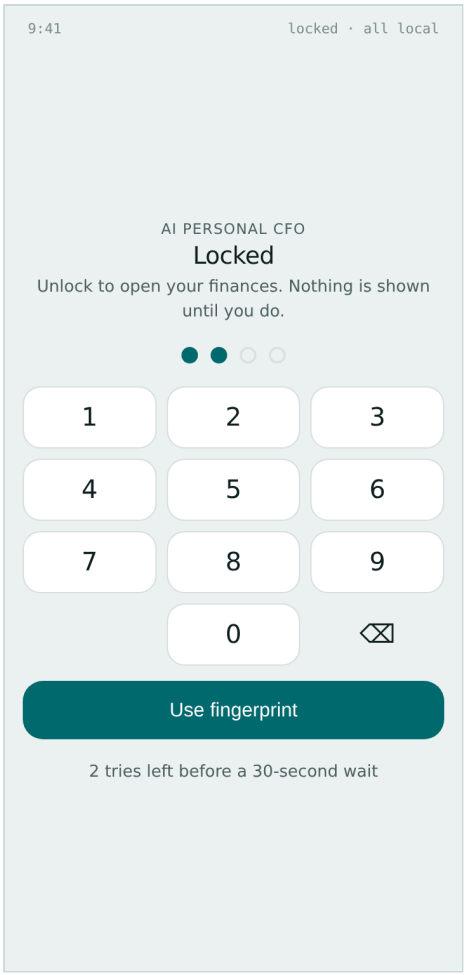
- A pledge with its exceptions stated is more credible than one without, and it is the same list the consents screen will show later — the user meets it twice and it agrees with itself (P-01, consistency).
- Two cards of equal weight made the sample path look like an equal choice. One filled button, one text alternative (move 5).
- “About 2 minutes” is the cheapest abandonment fix there is: people finish a setup whose end they can see.

2 · App lock

:app · SEC-002, §23.1 · moves 1, 5, 6



Today — the shipped render



Proposed — v2

What changed

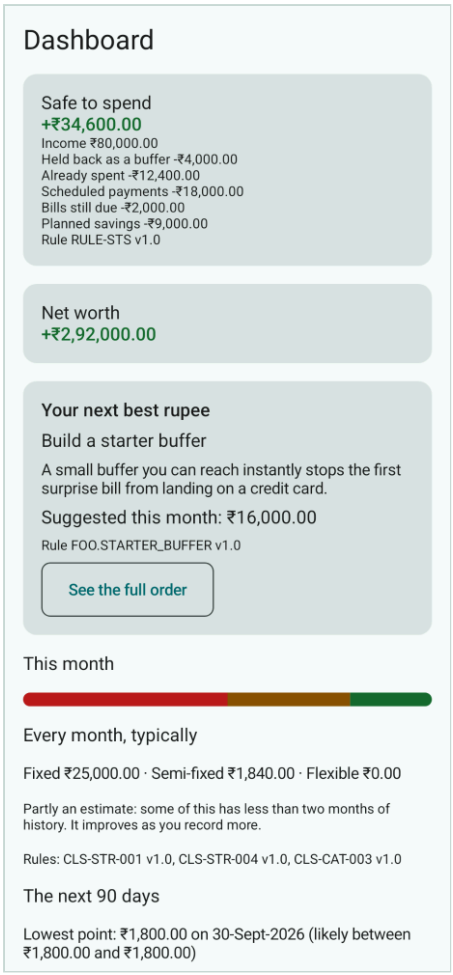
- The PIN text field becomes four dots and an on-screen keypad; fingerprint is the single filled action.
- Attempts remaining are stated under the keypad before a lockout, not after it.
- Nothing else is on the screen — no balance, no name, no figure.

Why

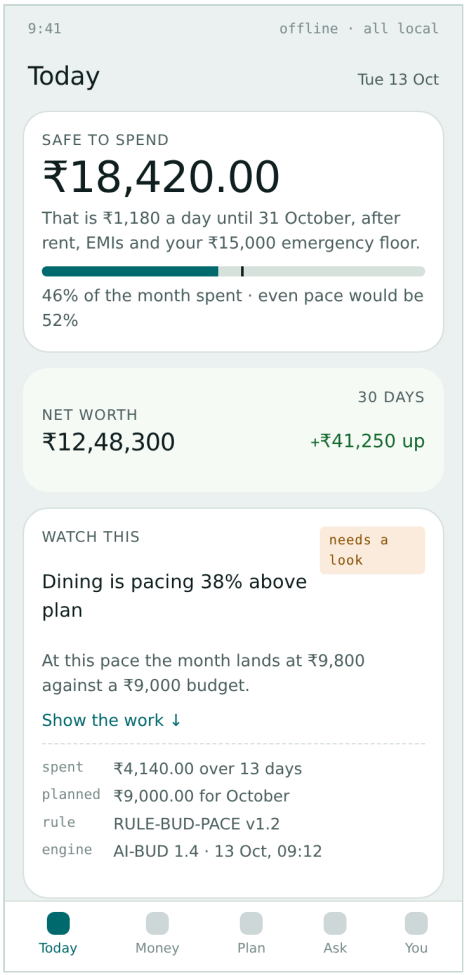
- A text field summons the system keyboard, which is what pushes the Unlock button below the fold at 200% font — a risk the screen's own source comments on. A fixed keypad cannot do that, and its keys are 100×48dp against WCAG 2.5.8's 44×44 minimum.
- Biometric first because it is the fastest path and the one SEC-002 prefers; the PIN stays visible as the fallback rather than hidden behind a link.
- “2 tries left” is visibility of system status (Nielsen #1) at the only moment it is actionable.

3 · Today — the dashboard

:feature:dashboard · §5.2, §14, AI-STS · moves 1, 2, 3, 4



Today — the shipped render



Proposed — v2

What changed

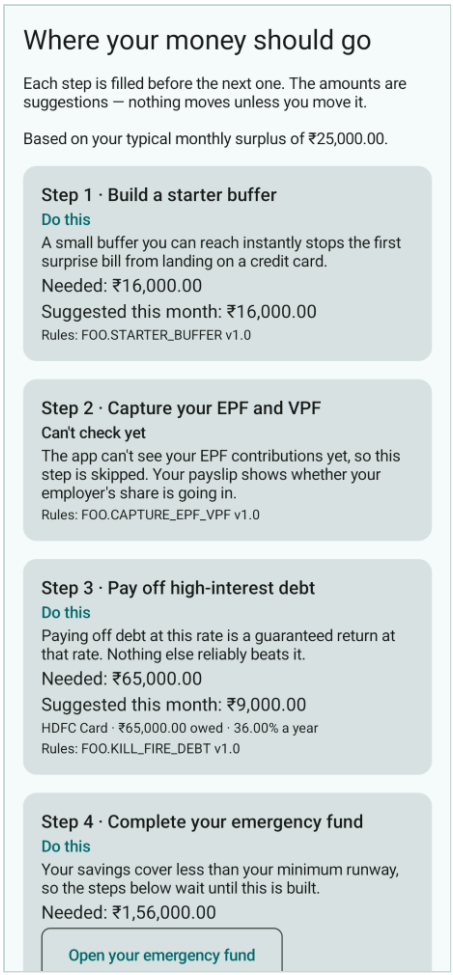
- Safe-to-Spend keeps the big type; the six amounts it was derived from move into *Show the work*.
- A meaning line turns the figure into an answer: what it is per day, until when, after what.
- The unlabelled stacked bar becomes a pace bar with an even-pace mark and a sentence saying where the month lands.
- The fourteen navigation rows become five tabs; what remains on Today is the one insight worth acting on and the next best rupee.

Why

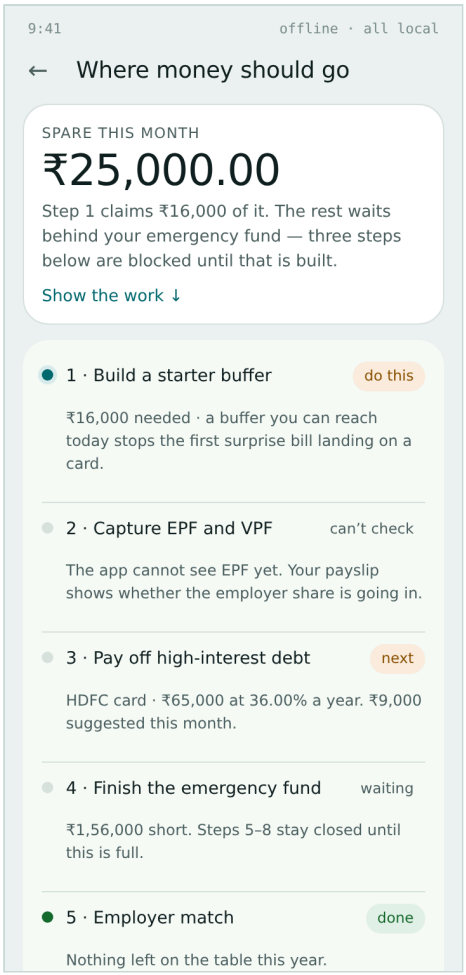
- Six ledger lines under the headline figure are evidence, not content. P-02 asks for the working to be available, not permanently unfolded — one tap, with the rule id and version in it (move 3, AI-ARC-003).
- A bar with no scale, no labels and no legend says less than the sentence beside it, and a bar that means something only by colour fails WCAG 1.4.1. The mark gives it a reference and the words carry the meaning.
- A home screen that is also the app's directory cannot lead with an answer. Material 3's navigation bar takes three to five destinations; this is five.

4 · Where money should go

:feature:dashboard · §36, FOO-002 · moves 1, 3, 5



Today — the shipped render



Proposed — v2

What changed

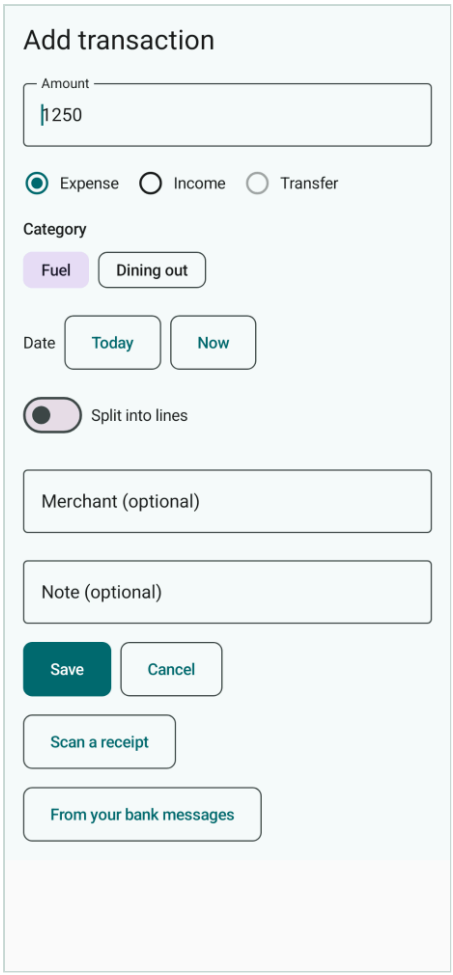
- The spare amount leads, and the meaning line says what step 1 claims of it and what is blocked behind the emergency fund.
- Eight stacked cards become a stage list: a status dot, the step, a word pill, one sentence.
- `Rules: FOO.STARTER_BUFFER v1.0` and the rest move into the card's disclosure.

Why

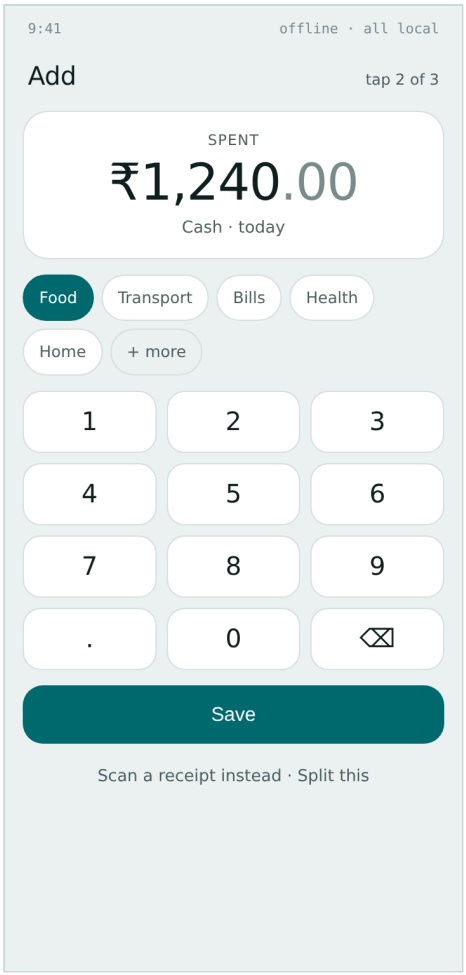
- §36's claim is that the *order* is the advice, so the sequence has to be visible at a glance. Eight equal cards lose it; a numbered stage list with a state per row keeps it.
- Status is a dot *and* a word — *do this*, *next*, *waiting*, *done* — so it survives greyscale and the privacy blur.
- “Amounts are suggestions. Nothing moves unless you move it” stays on the screen, because P-07 is the whole reason this list is allowed to rank a user's money.

5 · Add a transaction

:feature:transactions · FR-TXN-001/002 · moves 1, 5



Today — the shipped render



Proposed — v2

What changed

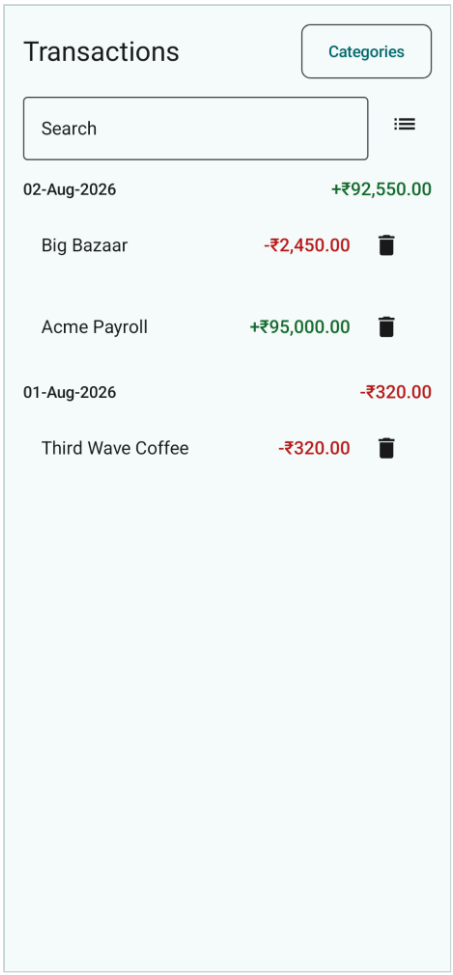
- The amount is the screen, with its own keypad; the system keyboard never opens.
- Expense / Income / Transfer becomes a segmented control; categories are 44dp chips with a + *more*.
- The date and account read as values (*Cash · today*) rather than as two buttons labelled *Today* and *Now*.
- One filled **Save**; the receipt and bank-alert routes become one text line beneath it.

Why

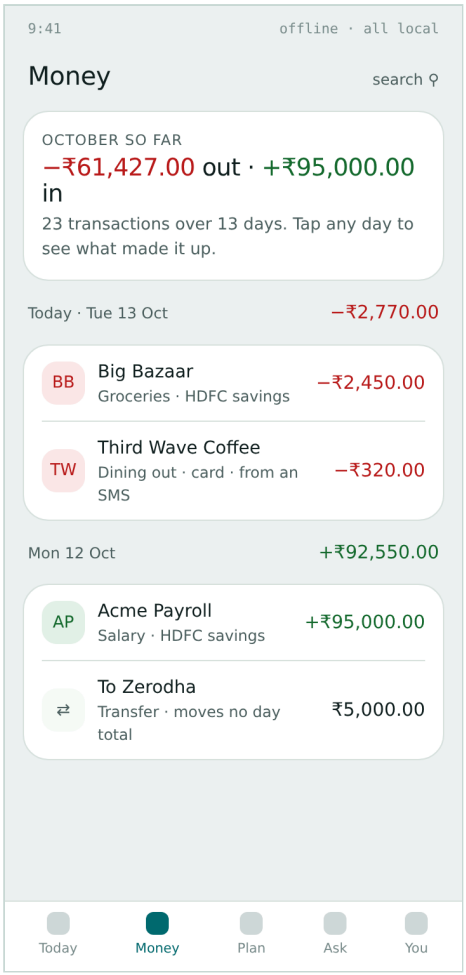
- FR-TXN-002 is a layout budget — one tap to reach, three to finish. A built-in keypad removes the focus-then-keyboard dance and makes the target area the biggest thing on screen.
- Radio buttons cost a row each and read as a questionnaire; a segmented control is Material 3's pattern for exactly-one-of-two-or-three opposites, and it is one tap in the same place.
- Two buttons labelled *Today* and *Now* state an action but hide the value, so the user cannot check what they are about to save without tapping (Nielsen #1).
- Four buttons at four widths at the end of a form is a choice the user did not ask for. One primary; the alternatives are routes, so they read as text.

6 · Transactions

:feature:transactions · FR-TXN-007 · moves 1, 4, 5



Today — the shipped render



Proposed — v2

What changed

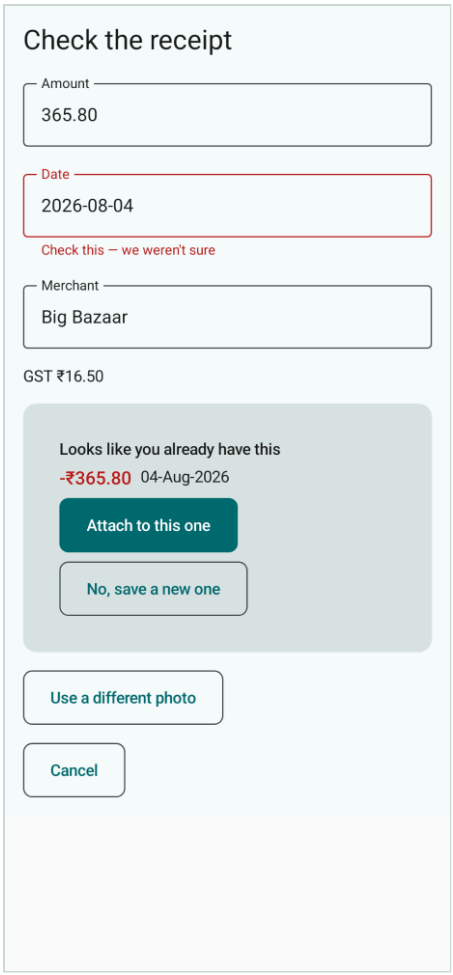
- A summary card answers the month first: what went out, what came in, over how many transactions.
- Day headers carry the day's net total; rows gain a leading glyph and a category·account note.
- The per-row delete icon goes; deleting happens inside the row's own screen, with the undo that already exists.
- Search moves to an app-bar action instead of a permanent field above the list.

Why

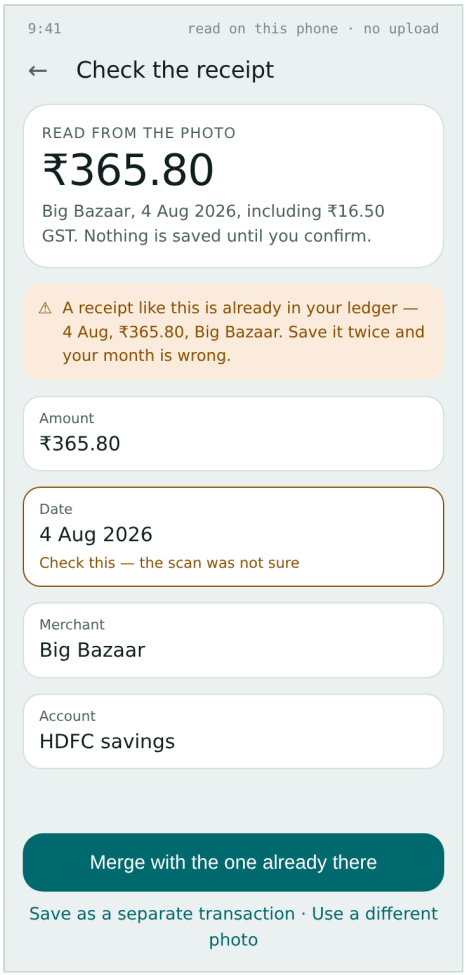
- A destructive action repeated on every row of a scrolling list is an accident waiting for a thumb. Material 3 keeps one primary target per list row and puts the rest in the destination.
- A ledger is scanned by shape before it is read; a glyph and a tone let the eye find the salary row without reading four lines (recognition over recall, Nielsen #6).
- The day total was already the thing FR-TXN-007 asks for — it just needed to look like structure rather than another row.

7 · Receipt review

:feature:transactions · FR-OCR-001…006 · moves 1, 5



Today — the shipped render



Proposed — v2

What changed

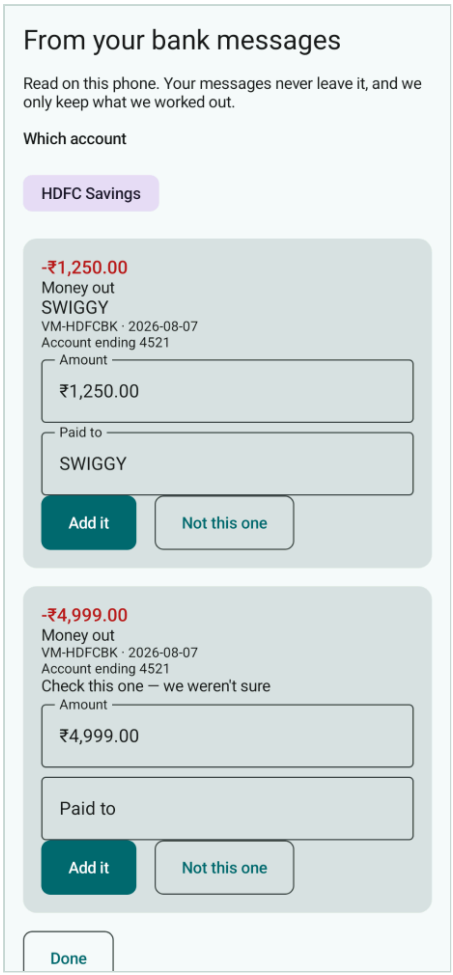
- The total read from the photo is the one answer; the meaning line carries the merchant, the date and the GST.
- The duplicate match becomes a banner above the fields, and **Merge** becomes the primary action.
- The field the scan was unsure of keeps its warning tone and gains a sentence saying so.
- Three stacked buttons become one primary and a text row.

Why

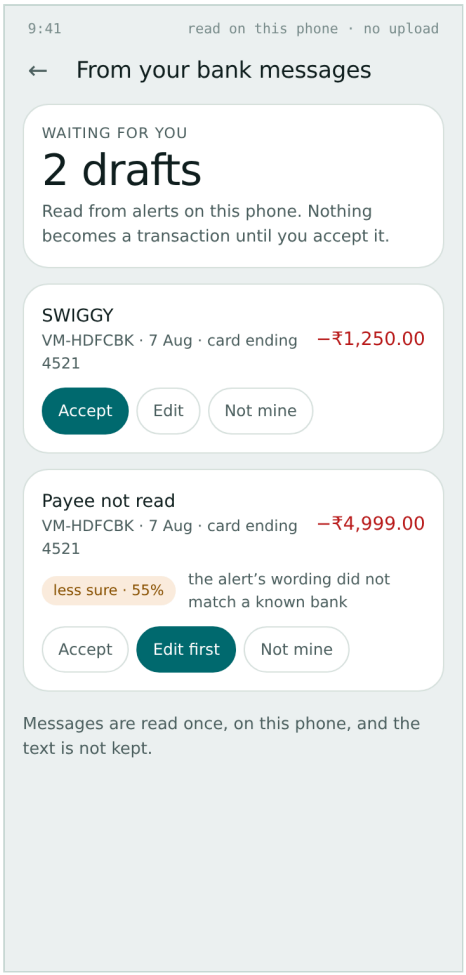
- FR-OCR-006's duplicate is the most consequential thing on the screen — it is the difference between one transaction and two — so it ranks above the fields it concerns.
- Merging is the safer default, and defaulting to the safe option is error prevention (Nielsen #5). Saving anyway stays available, as text.
- A red outline alone would be a colour-only signal; the sentence under the field is what makes the flag survive greyscale (WCAG 1.4.1).

8 · Bank-alert drafts

:feature:transactions · §18, §23 · moves 1, 5, 6



Today — the shipped render



Proposed — v2

What changed

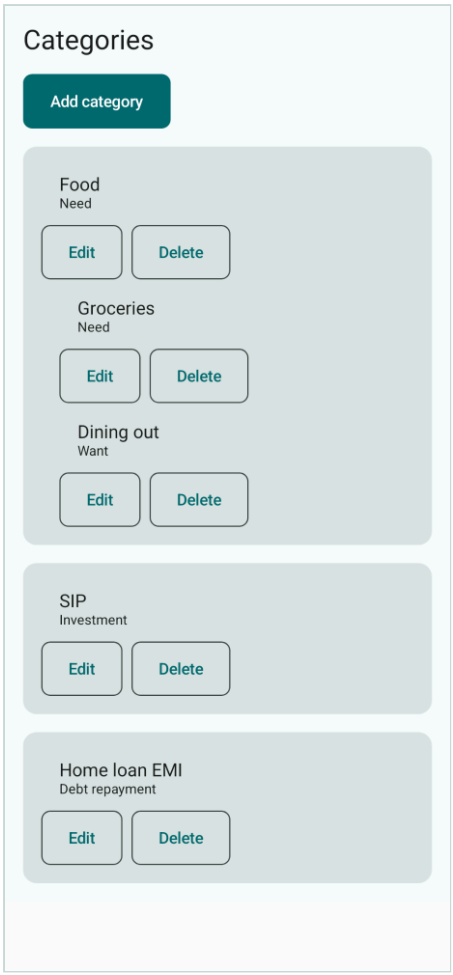
- "2 drafts" is the answer, with the standing line that nothing becomes a transaction until accepted.
- Each draft is a card with its amount, its source alert and its own Accept / Edit / Not mine.
- The uncertain draft states its confidence as a word and a number, and says what the parser could not read.
- The status line carries *read on this phone · no upload*; the footer says messages are read once and not kept.

Why

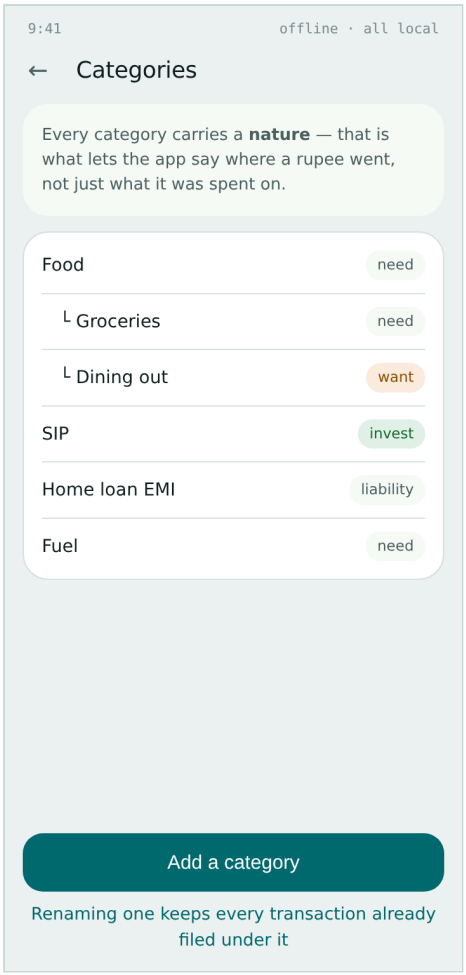
- P-07 is per draft, not per screen: a bulk *accept all* would turn a review into an import. Actions belong on the thing they act on.
- "Less sure · 55%" with the reason beside it lets a user judge the parse; a colour alone would ask them to trust it.
- This screen is the only place in the app that asks for `READ_SMS`, so the offer has to be legible *before* the system dialog appears (P-01).

9 · Categories

:feature:categories · FR-SET-001, §8.3 · moves 1, 5



Today — the shipped render



Proposed — v2

What changed

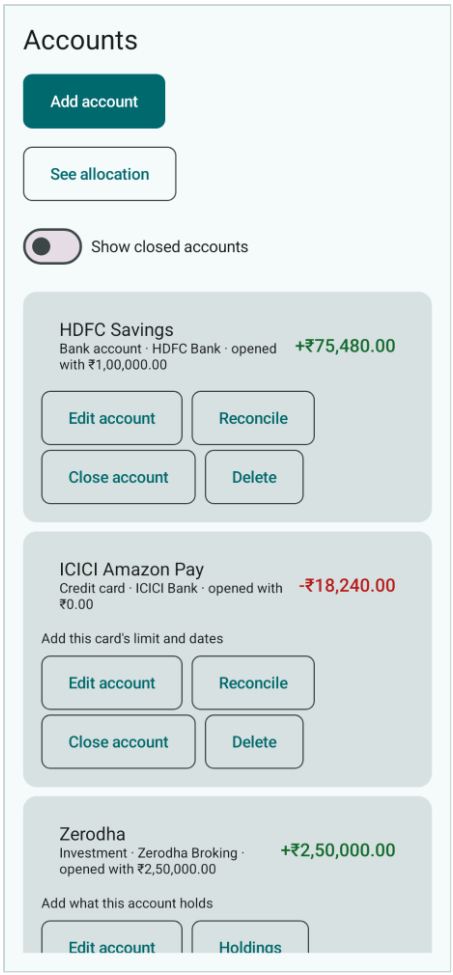
- Every row carries its nature as a pill — need, want, invest, liability.
- The hierarchy reads by indent; the editor stays a sheet over the list.
- One filled **Add a category**, and a line saying a rename keeps the history already filed under it.

Why

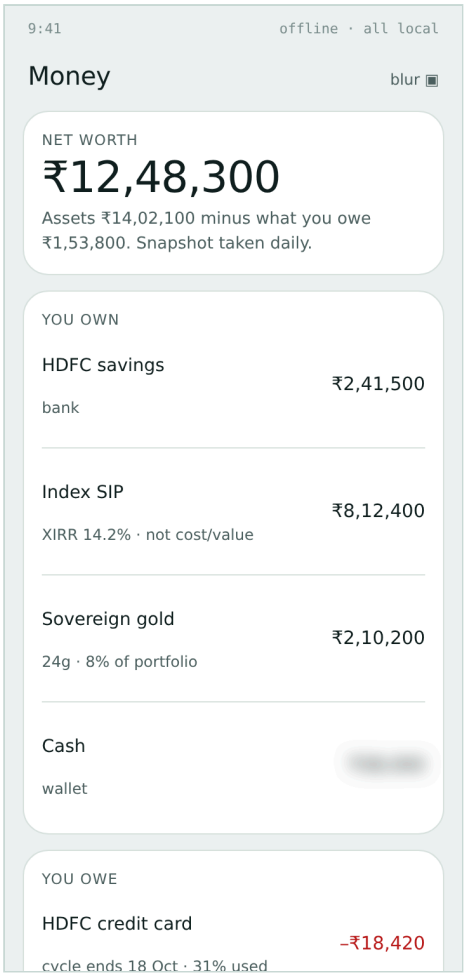
- Nature is what makes §8.3 possible — “what did this rupee become” — and it was invisible on the screen that owns it. A pill is a word, so it is greyscale-safe.
- The rename note answers the question this screen reliably raises, at the moment it is raised, instead of in help nobody opens (Nielsen #10).

10 · Money — accounts

:feature:accounts · FR-ACC-001/005/007 · moves 1, 2, 5



Today — the shipped render



Proposed — v2

What changed

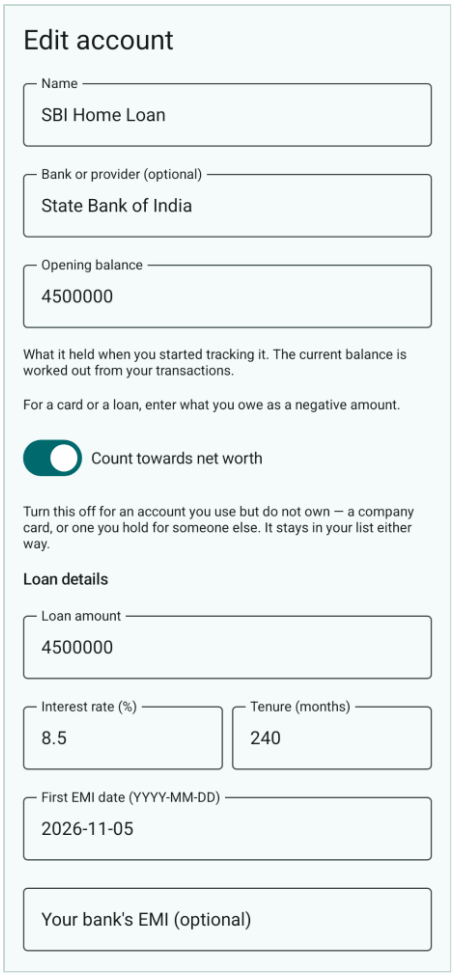
- Net worth becomes the one answer, with assets and liabilities named in the meaning line.
- Accounts group under **You own** and **You owe**; each row is one tap into that account.
- The four buttons per row — Edit, Reconcile, Close, Delete — move inside the account.
- A card states its cycle and utilisation in words; a loan states how much of the EMI is interest.

Why

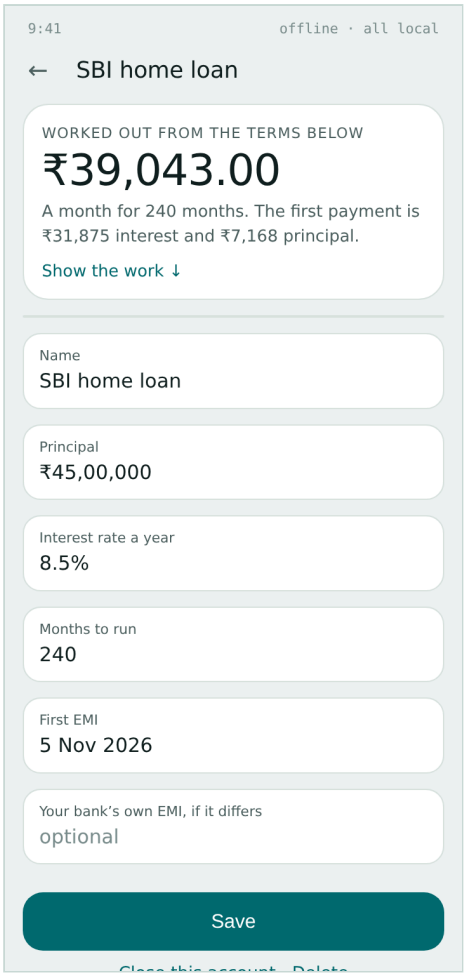
- Four full-width buttons on every row is the densest thing in the app and the clearest Material 3 violation: a list row has one primary target, and its secondary actions live in the destination or behind an overflow.
- “You own” / “You owe” is FR-ACC-005's own arithmetic — assets minus liabilities — made visible, so the net-worth figure above is checkable by reading down.
- A derived balance is the row's job (DB-001), so it is the only amount on the row; the opening balance moves inside.

11 · Account editor

:feature:accounts · FR-ACC-002/003 · moves 1, 2, 5



Today — the shipped render



Proposed — v2

What changed

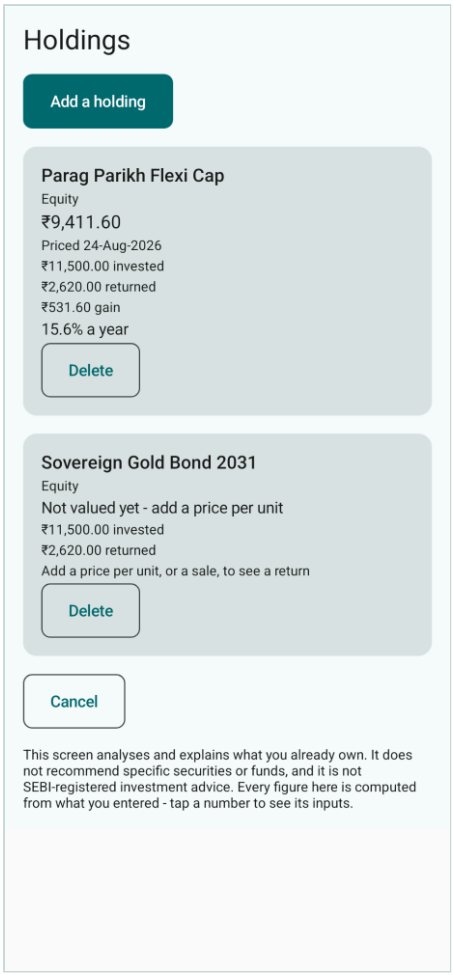
- The EMI the terms produce is the screen's one answer, with its first-payment split beneath it.
- The type picker becomes a segmented control; the type-specific section follows it.
- Fields show their values at a readable size; the rate is in percent, as typed.
- One filled **Save**; closing and deleting become text.

Why

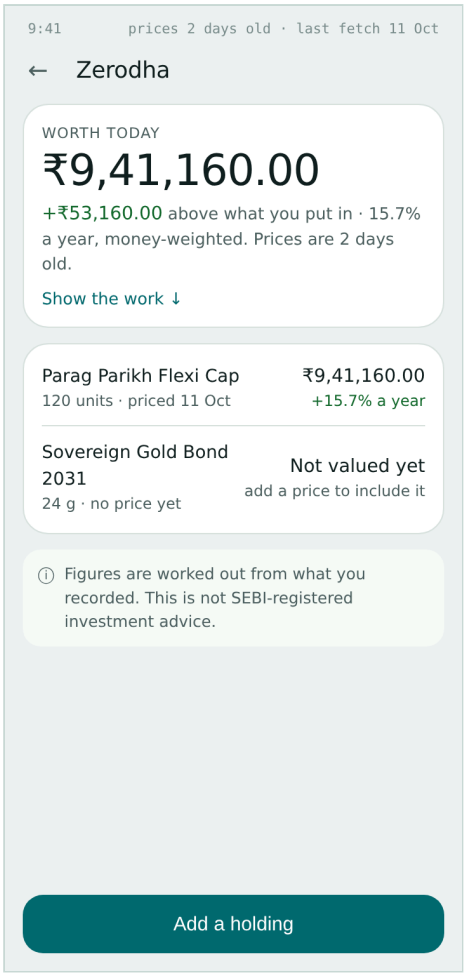
- A form whose output is a schedule should show the schedule. It turns data entry into a feedback loop — the user can tell a wrong rate from a right one by looking at the EMI (P-02, Nielsen #1).
- Basis points are an implementation decision (MNY-002), not something to make a borrower think about; the screen takes 8.5 and stores 850.
- Delete at the same weight as Save is how an account disappears by accident.

12 · Holdings

:feature:accounts · §11, AI-INV · moves 1, 2, 3, 6



Today — the shipped render



Proposed — v2

What changed

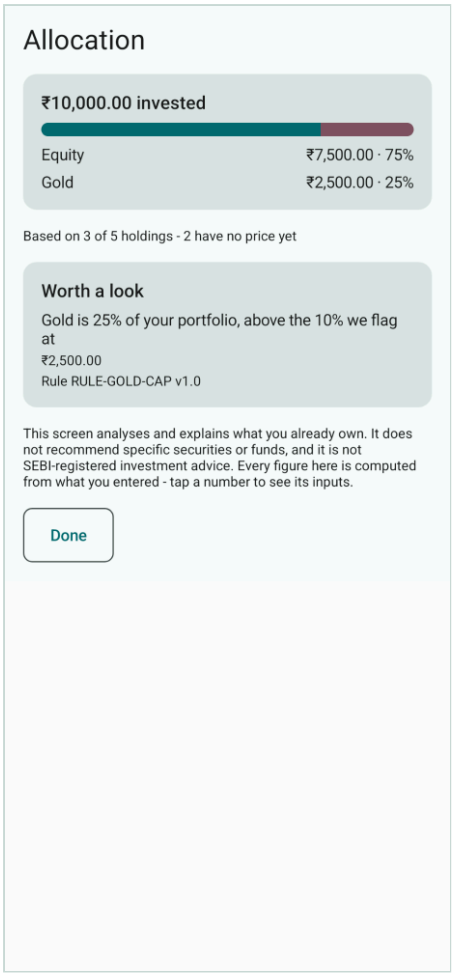
- The account's value leads; gain and money-weighted return sit in the meaning line.
- Price age moves into the status line, where it qualifies every figure below it.
- An unpriced holding says *Not valued yet* and offers the action that fixes it.
- §11.1's disclaimer becomes a banner at the foot rather than a paragraph in the content column.

Why

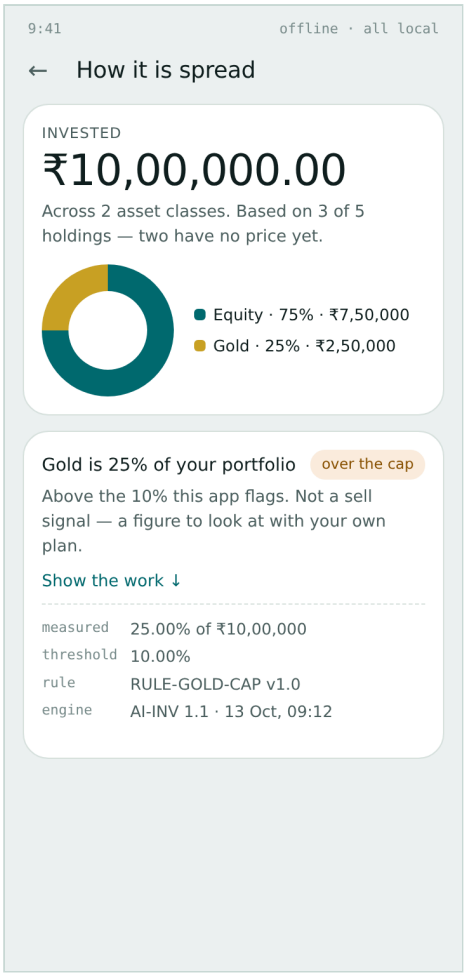
- A return without its basis is a claim. “15.7% a year, money-weighted” beside “+₹53,160 above what you put in” is the same figure said twice, which is what makes it checkable (P-02).
- Staleness applies to the whole screen, so it belongs in the chrome — one statement instead of a label repeated per row (move 6).
- “Not valued yet” is P-03 made visible: absence is a prompt, never a zero, and the prompt carries its own fix.

13 · Allocation

:feature:accounts · FR-INV-002, §11.2 · moves 1, 3, 4



Today — the shipped render



Proposed — v2

What changed

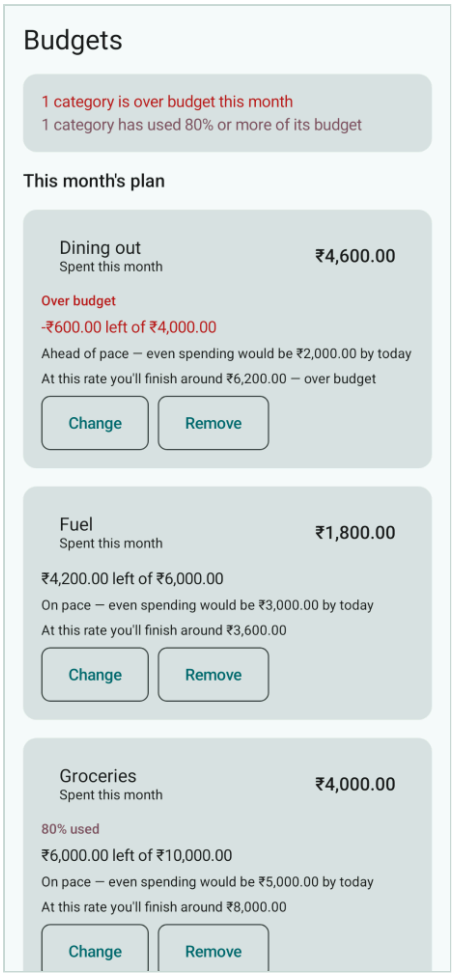
- A donut with a legend replaces the text split; every slice carries its share *and* its amount.
- The concentration flag becomes its own card with a word pill, and the rule id moves into its disclosure.
- Coverage — “based on 3 of 5 holdings” — moves into the meaning line under the total.

Why

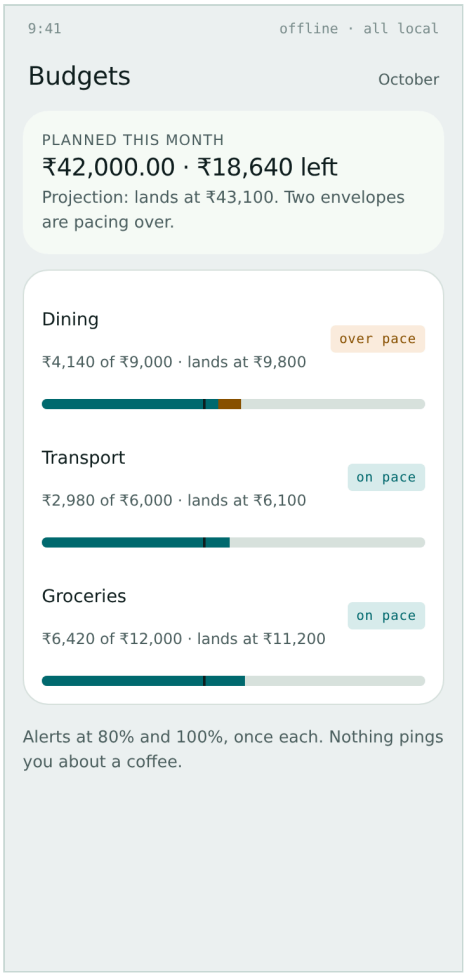
- The screen's one answer is a proportion, and a proportion is the one thing a bar cannot say at a glance. The legend carries the numbers, so the chart is never the only source (WCAG 1.4.1).
- “Not a sell signal” is on the flag itself, because a cap breach on a finance screen is read as an instruction unless it says otherwise (P-07).
- Coverage qualifies the whole split; printing it as a separate line let it be read as an aside.

14 · Budgets

:feature:budgets · FR-BUD-001…004 · moves 1, 4, 5



Today — the shipped render



Proposed — v2

What changed

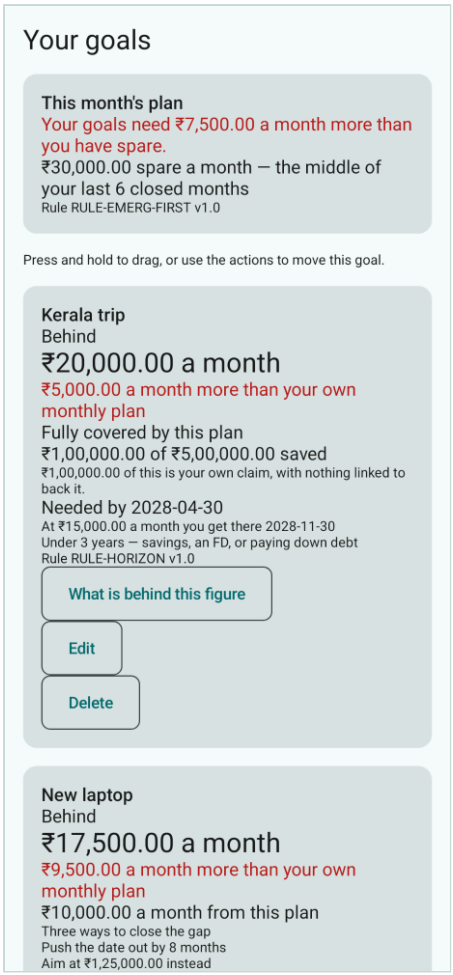
- The month's plan leads: planned, left, and where the month is projected to land.
- Each envelope is one row and one bar — spent, the even-pace mark, the overshoot in the warning tone.
- “Over pace” is a word beside the bar; **Change** and **Remove** move inside the envelope.
- The alert policy is stated once at the foot: 80% and 100%, once each.

Why

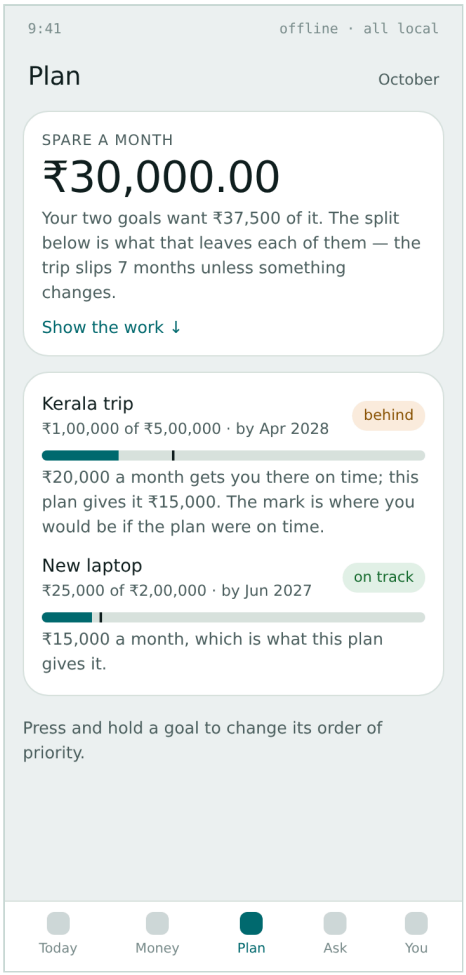
- Move 4 exists for this screen. One bar carries three facts that previously took three sentences, and the projection stays in words because “lands at ₹9,800 of ₹9,000” is the part a person acts on.
- Two buttons on every envelope is the accounts problem again: n × 2 equal actions, none of them the thing the user came for.
- Saying when it will interrupt you is how a budget feature earns the right to send notifications at all (§17.2, Nielsen #1).

15 · Goals

:feature:goals · §15, AI-GOAL, §36 · moves 1, 2, 4



Today — the shipped render



Proposed — v2

What changed

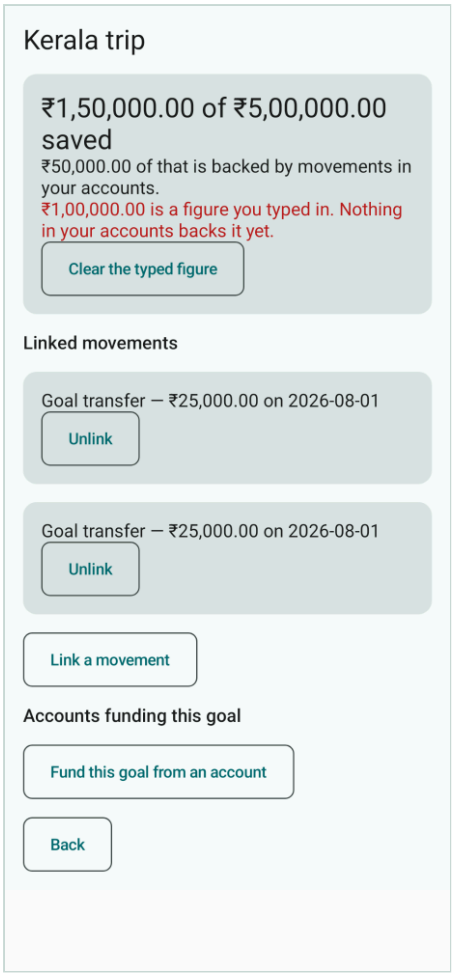
- The spare per month leads, and the meaning line says what the goals want of it and what that costs in time.
- Each goal is a row, a progress bar with a pace mark, and a pill — *behind* or *on track*.
- Ten lines of prose per goal become two: the required monthly against what the plan gives it.
- Three ragged buttons per goal go; the goal opens.

Why

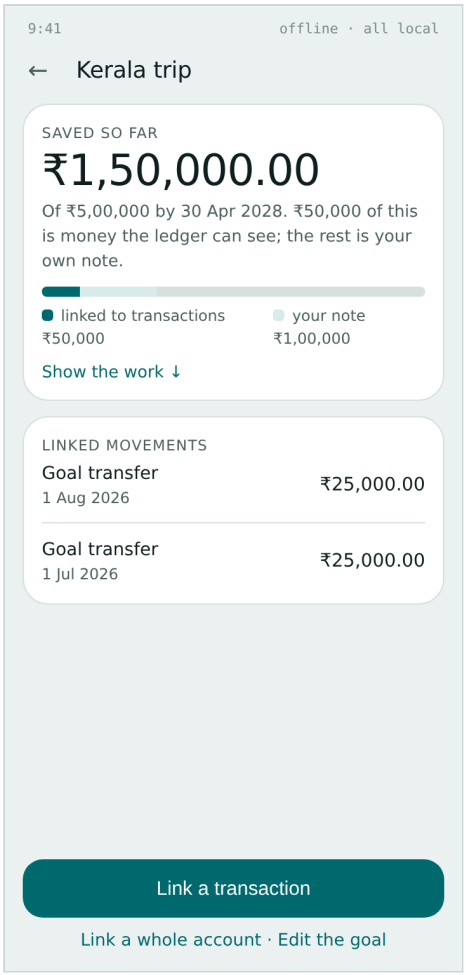
- Red was doing the wrong job. “Needs more than you have spare” is a caution, not an error — it moves to the warning role with a word, so the error role keeps meaning *stop* (P-02, WCAG 1.4.1).
- A goal is the textbook progress bar: a known target, a known position. Describing it in rupees twice was the longest way to say it.
- Dates read as *Apr 2028*, not `2028-04-30` — an ISO string is a storage format (TIM-002), not a sentence.

16 · One goal

:feature:goals · FR-GOAL-002/004 · moves 1, 3, 4



Today — the shipped render



Proposed — v2

What changed

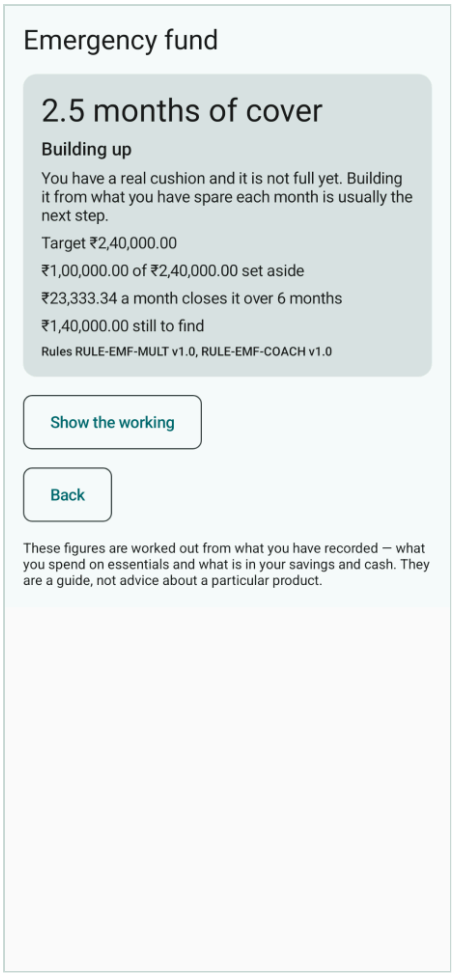
- Saved-of-target leads; the meaning line names the split between what the ledger can see and what is a note.
- One bar carries both parts, with a two-word legend naming each.
- The linked movements are listed under it; linking is the primary action.

Why

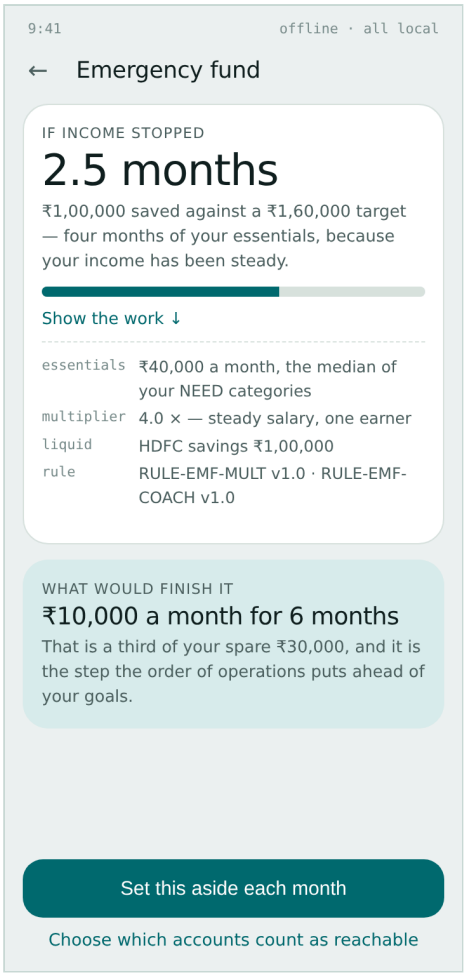
- Issue 7.4's subject is exactly this distinction, and it is a two-part quantity — the only honest form is one bar divided, with both parts named in words so it survives greyscale and the privacy blur.
- “Your note” rather than “unevidenced”: the user made that claim, and the design should not imply they were caught at something (P-07).

17 · Emergency fund

:feature:emergencyfund · §10.1, AI-EMF · moves 1, 2, 3



Today — the shipped render



Proposed — v2

What changed

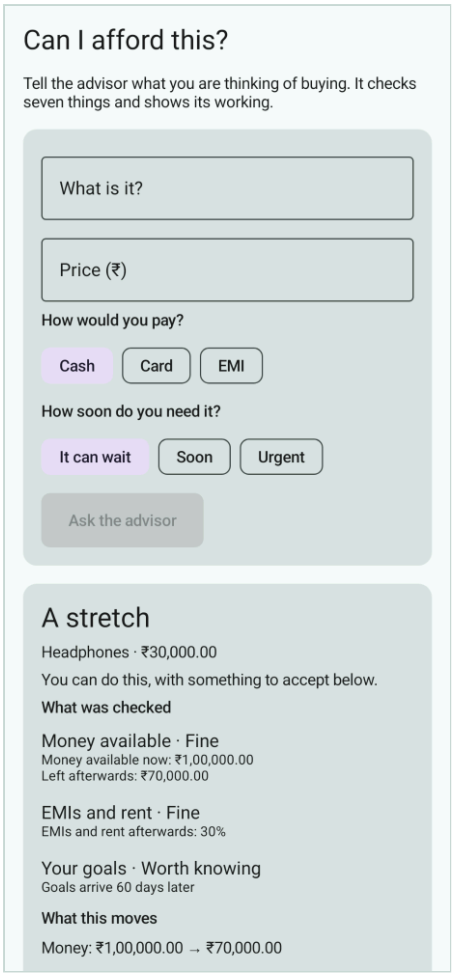
- The runway in months is the one answer — not the target amount.
- The amounts become the meaning line, and the multiplier explains itself there.
- §10.1's evidence is the inline disclosure: essentials, multiplier, which accounts counted, the rules.
- The coach line becomes its own card with the monthly figure and what it displaces.

Why

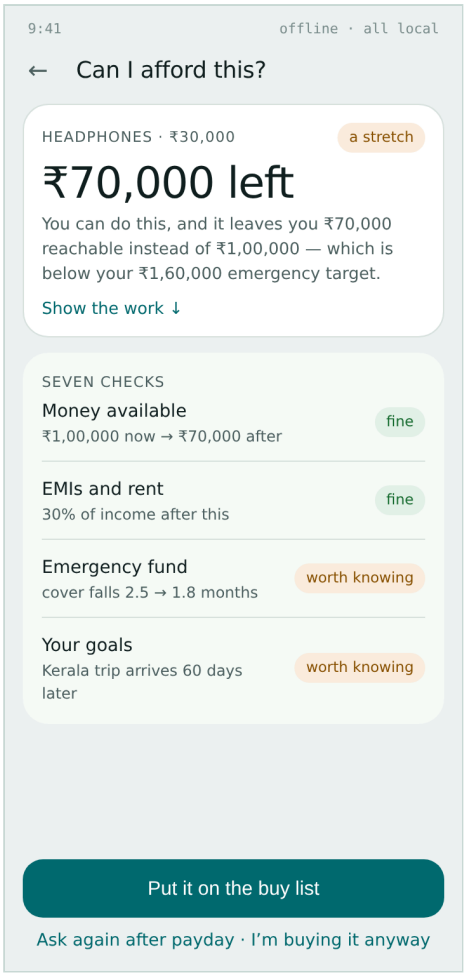
- “2.5 months” is what the user asked; “₹1,00,000 of ₹1,60,000” is how it was worked out. v1 led with the working (move 1).
- §10.1 asks every figure in the explanation to link to its evidence. One disclosure on the card is that link — and it is one tap, where a separate evidence screen would be one journey.
- The coach card names what the ₹10,000 competes with, because advice that ignores the other claims on the same rupee is advice that will be ignored.

18 · Purchase advisor

:feature:advisor · §13.2, FR-AI-003 · moves 1, 2, 3, 5



Today — the shipped render



Proposed — v2

What changed

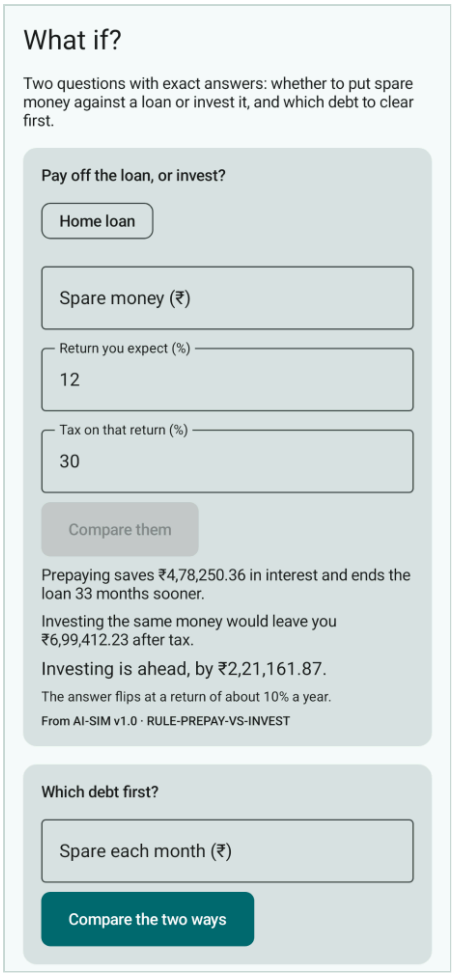
- The verdict becomes a pill beside the item and its price; the big figure is what the purchase *leaves*.
- The seven checks become rows: name, the figures that moved, a word pill.
- The primary action is **Put it on the buy list**; "I'm buying it anyway" stays as text.

Why

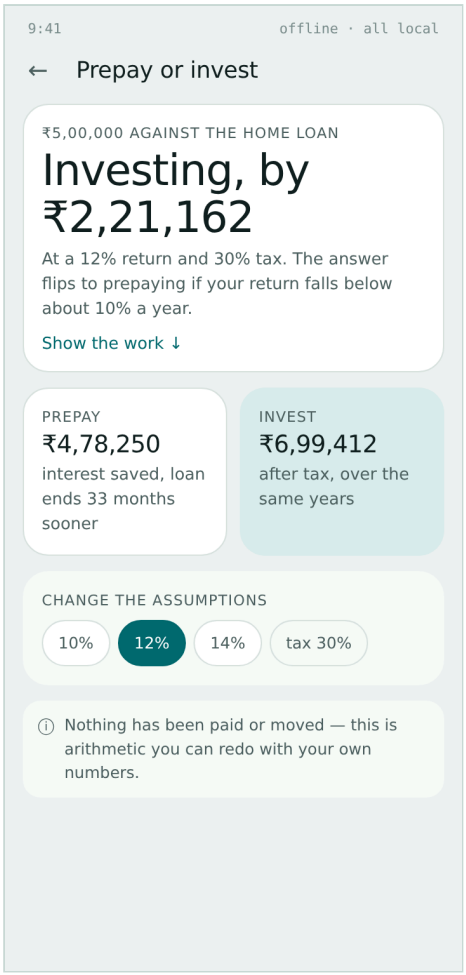
- The verdict is one word and was the largest thing on the screen; the figure a person actually decides on is what is left afterwards. Move 1 is about which figure, not about size for its own sake.
- Status as a pill rather than a chip row also removes the overflow that broke *Ur / ge / nt* across three lines at 200% font (ADR-0066) — the fix stays, and the pattern that caused it does not return.
- P-07 in one layout decision: the app's primary action is to *remember* the purchase, never to approve or refuse it, and the user's own decision is always on screen.

19 · Simulators

:feature:advisor · §36, §40.2, AI-SIM · moves 1, 2, 5



Today — the shipped render



Proposed — v2

What changed

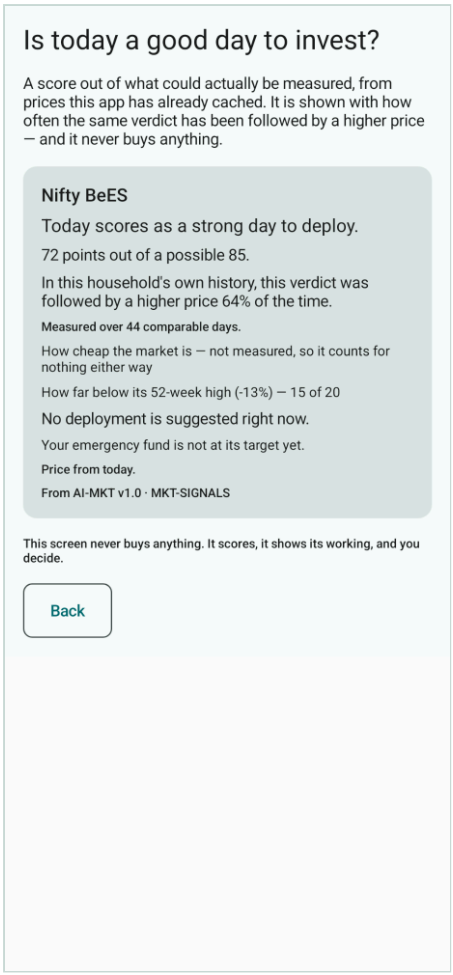
- The answer and its size lead — *Investing, by ₹2,21,162* — with the rate at which it flips.
- The two sides sit as two cards of equal size; only the one that wins is tinted.
- The assumptions become chips, editable without a keyboard.
- “Nothing has been paid or moved” becomes a banner.

Why

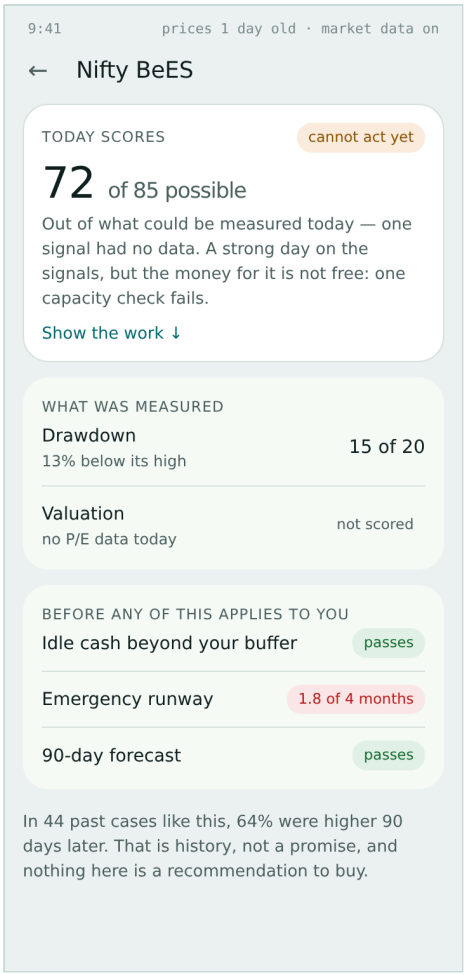
- Both sides visible, or the screen is a verdict and P-02 is broken. Equal cards say “these are comparable”; the tint says which one is ahead, and the sentence above says by how much.
- The flip rate is the most decision-useful number on the screen and it was the last line of a paragraph. A comparison that cannot be stress-tested is an opinion.
- This screen looks like a payment screen and is not one, which is exactly when a standing disclaimer earns its space (P-07).

20 · Opportunity

:feature:market · §30, AI-MKT · moves 1, 2, 6



Today — the shipped render



Proposed — v2

What changed

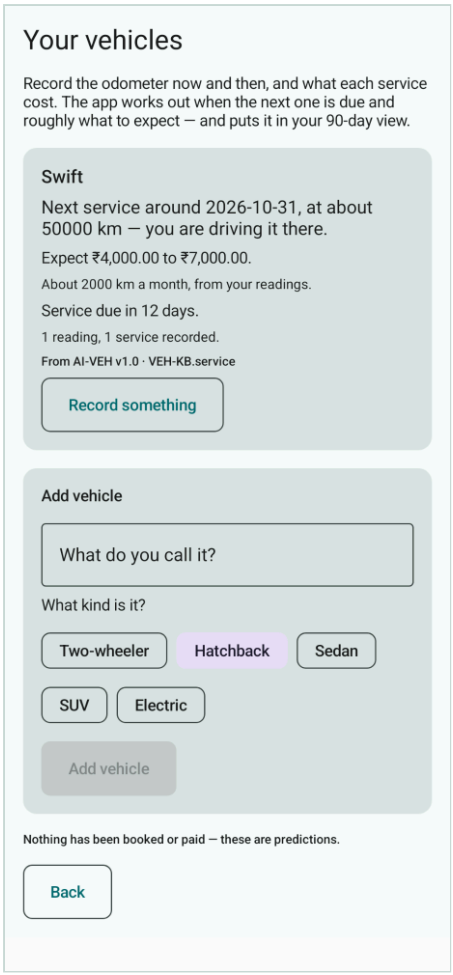
- The score is stated out of what was *possible* today, with a pill saying the money for it is not free.
- Signals and capacity gates become two short lists; an unscored signal says why.
- The hit rate is a sentence with its sample size and horizon.
- Price age sits in the status line.

Why

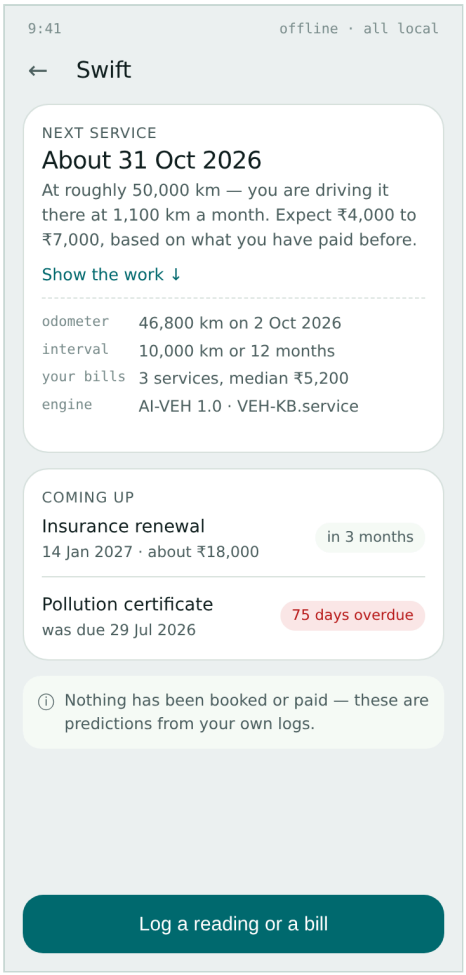
- This is the screen most likely to be read as a tip — its own source says so. So the failing capacity gate outranks the score visually: the answer is “good day, wrong time for you”, and both halves have to arrive together.
- “72 of 85 possible” refuses the false precision of a score out of 100 when a signal could not be evaluated (P-03).
- “In 44 past cases, 64% were higher 90 days later” is history with its denominator — a percentage without a sample is a sales pitch.

21 · Vehicles

:feature:vehicle · §12, AI-VEH · moves 1, 2, 3



Today — the shipped render



Proposed — v2

What changed

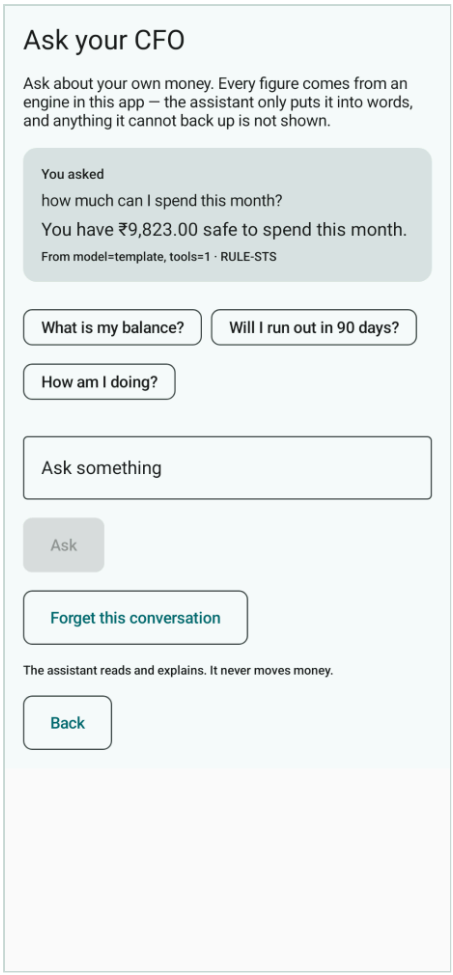
- The next service date leads, with the cost range and the reason for it in the meaning line.
- The evidence — odometer, interval, your own bills — becomes the disclosure.
- Renewals are rows with word pills; an overdue one says how many days.
- “Nothing has been booked or paid” becomes a banner.

Why

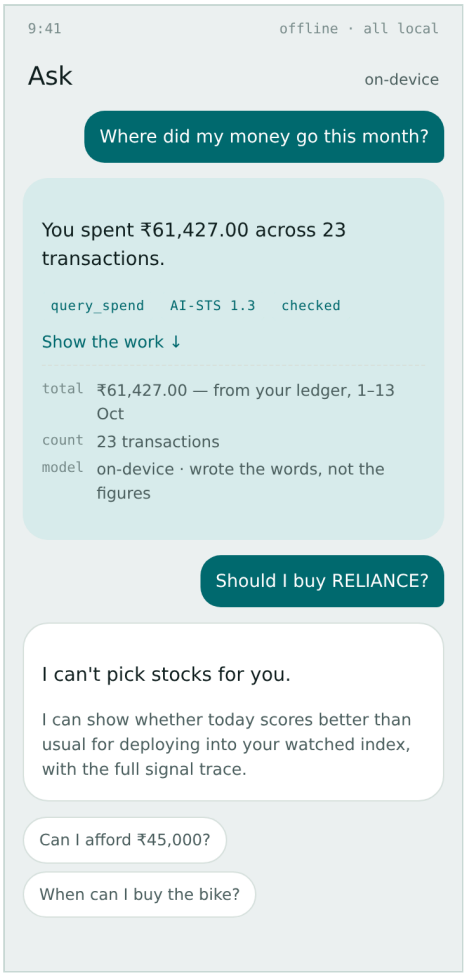
- A prediction without its basis is a guess, and this screen predicts money. Naming the three inputs is what makes “about ₹5,200” arguable (P-02).
- “75 days overdue” is a word and a number; the tone only reinforces it. An overdue certificate that signalled by colour alone would be invisible in the privacy blur.
- The range narrows from the knowledge base to the household's own bills, and the screen says which — that is the difference between a generic estimate and the user's own history.

22 · Ask — the assistant

:feature:chat · §19, AI-CHAT, AI-ARC-004 · moves 1, 3, 6



Today — the shipped render



Proposed — v2

What changed

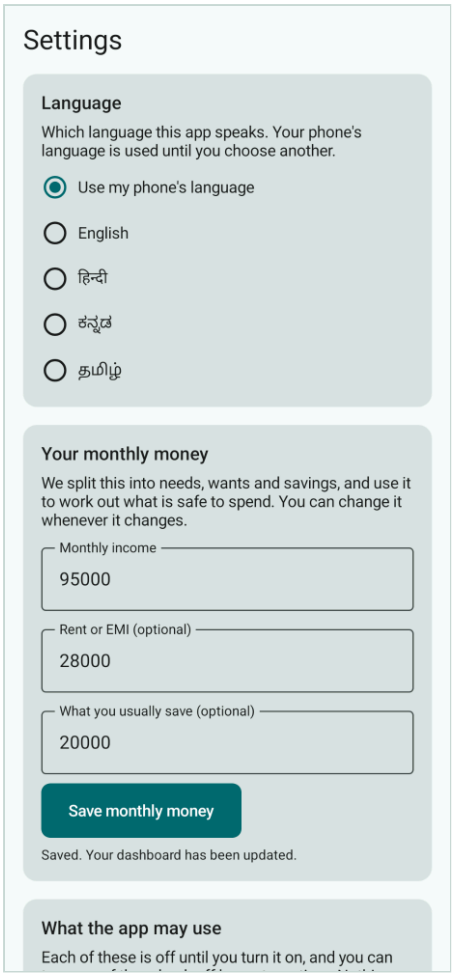
- It becomes a conversation: the question and the answer are separate bubbles with the composer pinned at the bottom.
- The tools used and the rule behind each figure sit on the answer as a compact line, with the full trace one tap away.
- Suggested questions sit above the composer, where a thumb can reach them.
- A refusal says what the assistant *can* do instead; **Back** becomes the app bar.

Why

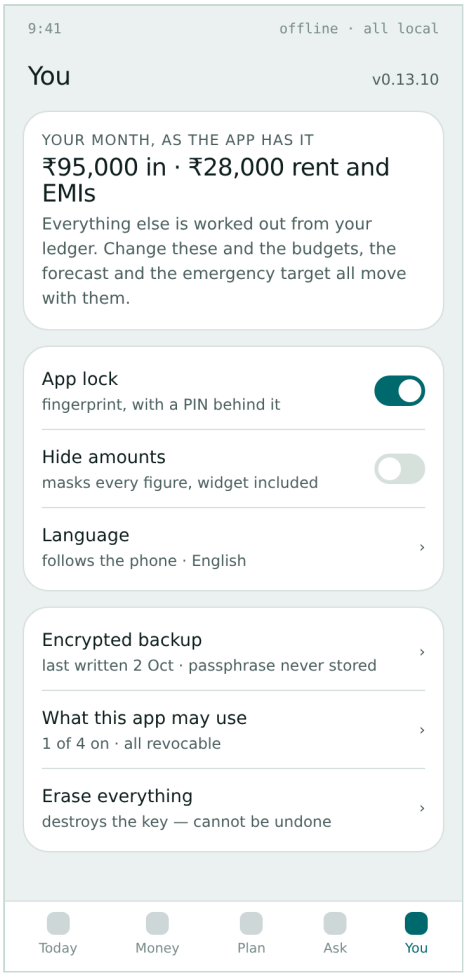
- A chat that does not look like a chat forfeits everything the pattern gives for free — who said what, in what order, with what still pending (Nielsen #6).
- The guardrail is this feature's whole safety case (AI-ARC-004, P-03), so what the answer is built from belongs on the answer, not in a footer: tool, engine version, and “the model wrote the words, not the figures”.
- A refusal that offers the adjacent capability is a redirect; one that only says no is a dead end (Nielsen #9).

23 · You — settings

:feature:settings · FR-SET-001, SEC-002/005 · moves 1, 5, 6



Today — the shipped render



Proposed — v2

What changed

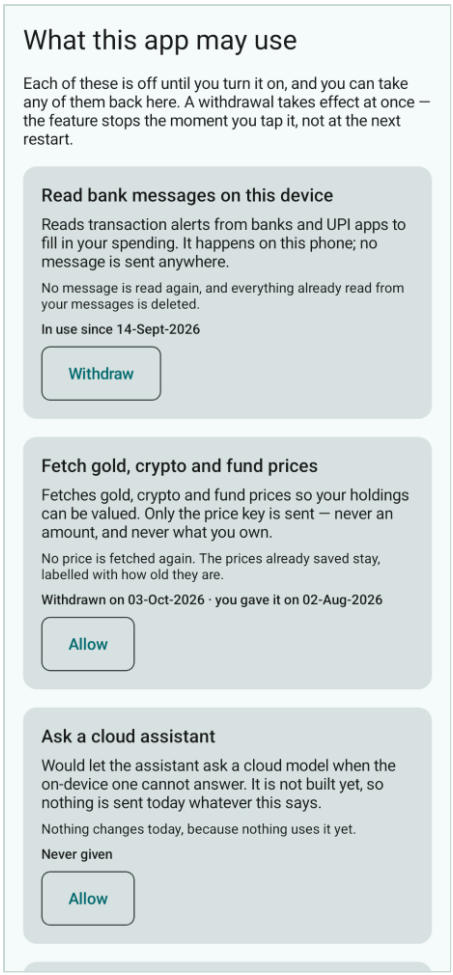
- The money plan leads, because everything else on the screen is derived from it.
- Each switch gains a line saying what it changes, not what it is called.
- Backup, consents and erase become rows into their own screens, each showing its current state.
- The encrypted-backup warning moves to the backup screen, beside the button it qualifies.

Why

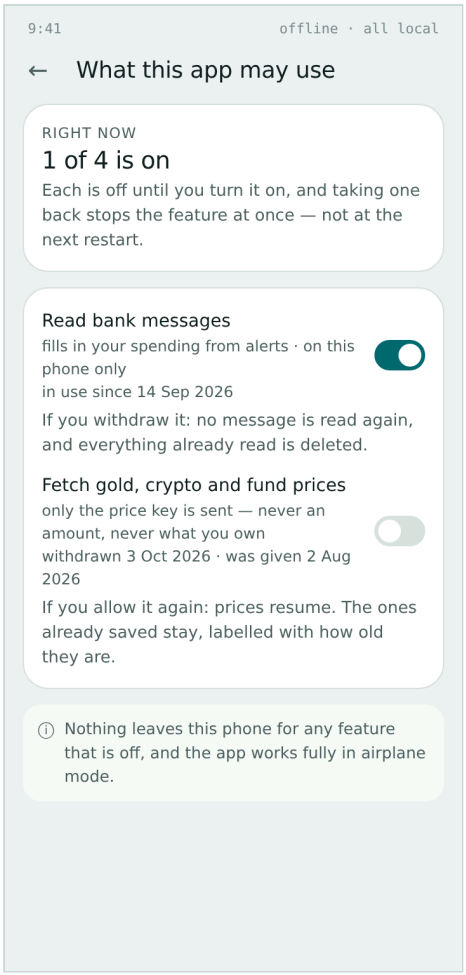
- The shipped screen is a form, two toggles and an encrypted-backup flow in one scroll; the consequential parts are in the middle of it. Grouping by *what this changes* makes the consequences legible.
- SEC-005's irrecoverability sentence has to be where the irrecoverable action is taken. On a long settings scroll it is read past; on the backup screen it is unavoidable.
- “1 of 4 on” on the consents row is the state, not the label — the user can audit from the list without opening it (move 6).

24 · Consents

:feature:settings · §23, P-01, DPDP · moves 1, 6



Today — the shipped render



Proposed — v2

What changed

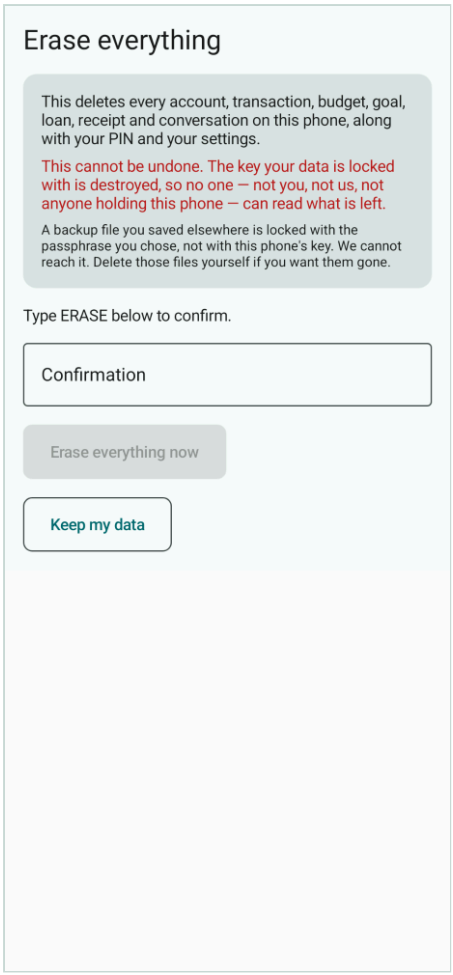
- “1 of 4 is on” leads, with the promise that withdrawal takes effect at once.
- Each row keeps its purpose, both dates, and — directly beneath it — what happens if it is withdrawn.
- The standing promise that nothing leaves the phone becomes a banner.

Why

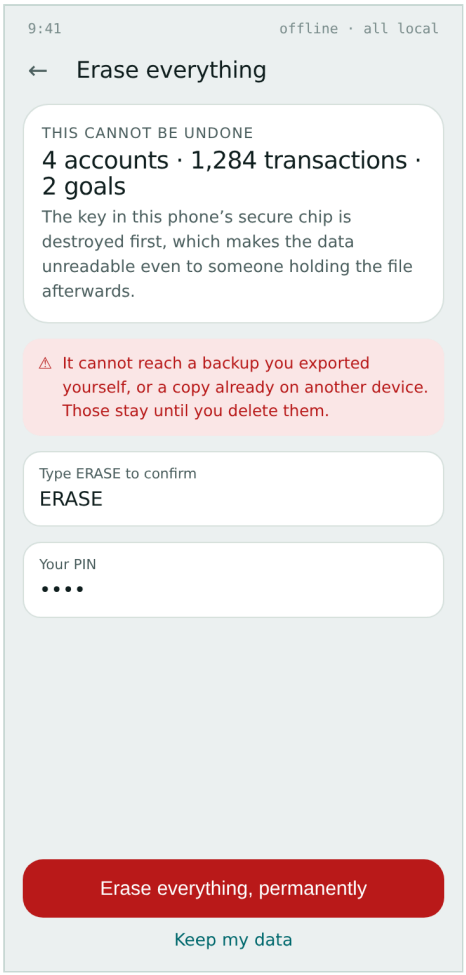
- The shipped screen is already the best-written in the app; what it lacked was a shape. The row order — what it does, since when, what stops — is the order a person asks the questions in.
- Keeping both the grant and the revoke date is what makes this an audit trail rather than a switchboard (P-01, DPDP), so neither is dropped for tidiness.
- A withdrawn consent keeps its row and its history instead of resetting to a blank switch: that record is the evidence the promise was kept.

25 · Erase everything

:feature:settings · §23, §34, SEC-003 · moves 1, 5



Today — the shipped render



Proposed — v2

What changed

- What will be destroyed is counted — accounts, transactions, goals — and that count is the one answer.
- The caveat about your own exported backups becomes a danger banner.
- The confirmation word and the PIN are fields on the same screen, not two steps.
- The destructive action is the only filled button in the app drawn in the negative role, and it is labelled with the verb and its finality.

Why

- A destructive confirmation has to state its scope. “Everything” is not a scope; “4 accounts · 1,284 transactions · 2 goals” is (Nielsen #5).
- The one thing a user cannot discover afterwards is what the erase could *not* reach, so it is the loudest thing before it.
- The negative role is reserved for this screen. A button that means *stop* everywhere else keeps its meaning because it is never used decoratively (P-02's colour rule, applied to actions).

What this would cost to build

Nothing here needs a new engine, a new dependency, a new colour seed or a new font. The figures are the ones the app already computes; what moves is hierarchy, order and where the evidence lives. That keeps the work in `:feature:*` and `:core:designsystem`, where it is covered by the Paparazzi baselines that already exist.

Work	What it touches	Size
The five-tab shell	A navigation bar in `:app`'s nav graph with a stack per tab. Typed routes already exist (ARC-001); this adds the host and the back stacks.	One change, every screen benefits
Four new components	`:core:designsystem` — status pill, segmented control, stage row, banner. Each is a composition of existing tokens; none needs a new one.	Small, and reused 20+ times
The *Show the work* disclosure	One expandable card section fed by the `EngineProvenance` every result already carries (AI-ARC-003). Nine screens stop printing provenance inline.	Small, mechanical
Per-screen re-layout	Each `*Content` composable, reordered. The `UiState` shapes do not change — every figure in these mockups already exists in the state that screen holds.	25 screens, independent of each other
New strings	The meaning lines. Every one is a resource with ICU plurals, translated into Hindi, Kannada and Tamil like the rest (`CfoHardcodedUiString` fails the build otherwise).	The largest single cost
Baselines	Re-record light, dark and 200% for each screen touched, and look at every image before committing (ADR-0066).	Per screen, unavoidable

> **What would need an ADR.** Nothing in this document changes the seed, the font or a non-negotiable, so none of it needs one on those grounds. Two decisions do: **the five-tab navigation shell**, because it changes a runtime call path and `FLOW.md` with it; and **moving provenance from the content column into a disclosure**, because P-02 is a golden rule and "one tap" is a reading of it that should be written down rather than assumed.

What this design does not fix

- **It does not make the app faster.** The dashboard's work is the engines', and nothing here changes what they compute or when.
- **It does not settle the widget.** The Glance widget has its own constraints and is out of scope; it should follow once the tokens here are settled.
- **It does not design the error states.** Each screen has an `errorCode` path and a loading path; those deserve the same treatment as the empty states and did not get it here.
- **It proposes English only.** Every meaning line is longer than the label it replaces, and Hindi, Kannada and Tamil run longer again — the 200%-font baselines in those locales are where this design will first be proved wrong.

How to disagree with it

Each screen above names the move, the standard and the project rule behind its layout. A reviewer who thinks a screen is wrong can say which of the three is being misapplied — which is a shorter conversation than one about taste. The mockups are editable HTML in `docs/design/redesign/mockups/`; changing one and re-rendering takes a few seconds, and the stylesheet is the v2 token table verbatim, so a change that looks wrong there is wrong in the app too.