


## EVIDENCE-LED PRODUCT DECISIONS

# Product Discovery, Behavioural Science & Experimentation


I combine behavioural science, product discovery and experimentation to understand what is happening in a user journey, investigate why it is happening, turn that evidence into product decisions, and test what actually changes.

**CASE 01 · MRUNNERS**

## Can motivational support adapt to the runner, not just the run?

Four years taking a psychology-informed running intervention from behavioural discovery through product definition, Android delivery and real-world evaluation.


Discovery → Product decisions →  
Delivery → Experimentation

**CASE 02 · ROYAL REST FINDER**

## How do you reduce uncertainty when someone urgently needs a public restroom?

An independent product taken from a recurring real-world problem to a functioning private beta in approximately seven weeks.

Problem framing → MVP → User  
feedback → Technical verification

**CASE 03 · DOCTOLIB**

## What if the hardest healthcare decision happens before the booking funnel?

An independent product diagnosis exploring practitioner selection, referral and cost uncertainty for people unfamiliar with the French healthcare system.

Journey diagnosis → Concept →  
Prioritisation → Validation plan

**ACROSS THESE CASES**

Behavioural insight · Product discovery · Mixed-method research · Prioritisation · Experimentation · Cross-functional delivery

## IN A GLANCE


# Starting is only part of the challenge. Staying engaged is another.

People can decide they want to start running and still struggle to maintain the behaviour over time.

MRunners focused on that part of the journey: what happens after the intention already exists, and whether more personally relevant motivational support can help sustain engagement.

## Can motivational support adapt to the runner, not just the run?

We explored whether psychology could help personalise motivational support to how different people pursue goals, making that support feel more relevant and potentially helping them stay engaged with running.



Functioning MRUNNERS Android intervention, September 2026.

## WHAT WE BUILT



### What worked

**Tailored motivational support made parts of the experience feel more positive.**

Participants responded more positively when motivational messages better matched how they approached their goals.

### What remained unresolved

**A better experience did not clearly translate into behaviour change.**

People responded more positively to the tailored experience, but we could not show that this changed their running behaviour.

## WHAT THIS CASE DEMONSTRATES

Behavioural product strategy

Discovery-to-delivery ownership

Experimentation

Evidence-led product judgement

#### ROLE

Day-to-day product & research lead under academic supervision

#### DURATION

Oct 2022 - Sep 2026

#### DELIVERY


3 Android versions evaluated in lab and real-use settings

#### TEAM

8 operational contributors under 3 supervisors

# Supporting the runner across the week

The product translated behavioural evidence into concrete moments of support rather than treating personalisation as a cosmetic layer.



## 1. Turn intention into a concrete plan

Runners could set weekly goals by sessions, duration and distance, then decide when they intended to run.

**Why it mattered:** control over goals and schedules repeatedly emerged as important, and planning gave the system context for when support might be useful.



## 2. Support the runner at the right moment

Notifications were linked to planned sessions, progress and contextual conditions rather than operating as generic reminders.


**Why it mattered:** the hypothesis was not simply “send more motivation”. It was whether the type and timing of support could feel more relevant to the person receiving it.



## 3. Make progress visible

Users could record activity, compare it with personal goals and contribute anonymously to a shared weekly challenge.

**Why it mattered:** competence and social motivation meant different things to different runners, so the product combined individual progress with collaboration rather than only ranking and competition.



## CONTEXT ALSO MATTERED

### Weather · planning · activity history · weekly feedback

Contextual information influenced perceived readiness to run, while activity history and progress created a feedback loop across the week.

## Product insight

**Personalisation was not about cosmetic differences.** It was about changing the kind of support people received according to how they approached goals, progress, social motivation and running context.

# Research only mattered if it changed what we built

Three decisions show how user evidence, technical constraints and delivery risk became product choices rather than reports.

## 1 Garmin integration → manual activity entry

|                    |                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>PROBLEM</b>     | Automatic activity retrieval remained unreliable, the project was behind schedule, and downstream evaluation needed dependable data. |
| <b>DECISION</b>    | Stop allowing Garmin integration to block the product and move to manual activity entry on Android.                                  |
| <b>TRADE-OFF</b>   | More effort for users and potentially less complete activity data.                                                                   |
| <b>CONSEQUENCE</b> | A less elegant but controllable workflow allowed development and field evaluation to continue.                                       |
| <b>LESSON</b>      | <b>Reliability beat technical ambition.</b>                                                                                          |

## 2 A motivational product with broken notifications does not work

|                  |                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------|
| <b>PROBLEM</b>   | Notifications were central to the intervention, but the implemented logic was not reliably delivering them. |
| <b>DECISION</b>  | Reduce trigger complexity, correct the critical logic and systematically test every planned notification.   |
| <b>RESULT</b>    | <b>27/27 planned notifications verified across six account/device configurations.</b>                       |
| <b>TRADE-OFF</b> | The recovery took approximately three months and delayed the field study.                                   |
| <b>LESSON</b>    | <b>Protect the mechanism the product exists to deliver.</b>                                                 |

## 3 The collective challenge frustrated users

|                  |                                                                                                      |
|------------------|------------------------------------------------------------------------------------------------------|
| <b>SIGNAL</b>    | <b>10/10 initial testers</b> expressed frustration with the challenge calculation.                   |
| <b>DIAGNOSIS</b> | The target could change during the week as individual goals changed, making progress unpredictable.  |
| <b>DECISION</b>  | Recalculate the collective target once each week and keep it stable through the week.                |
| <b>FOLLOW-UP</b> | Seven participants reviewed the revised approach, with qualitatively better understanding/reception. |
| <b>LESSON</b>    | <b>A theoretically interesting feature still has to make sense to the user.</b>                      |

# A better experience is not automatically a behaviour change

Four-week field evaluation · 25 completers · tailored vs neutral motivational-message conditions

## WHAT WORKED MORE CLEARLY

### Experience

Tailored motivational messaging improved parts of the subjective experience, including positive emotional response and joy during app use.

People felt better about parts of the experience when the messages better matched how they approached their goals.

## WHAT REMAINED LESS CLEAR

### Behaviour

Behavioural and broader outcome differences remained inconclusive.

People responded more positively, but we could not show that this changed their running behaviour.

## What this changed in my product thinking

Personalisation looked most promising as an experience strategy. The next product question is whether that improved experience can produce reliable, sustained behavioural effects.

## WHAT I WOULD DO NEXT

### Instrument behaviour better

Record app visits, notification interactions and more complete activity behaviour.

### Automate activity capture carefully

Revisit automated sources without sacrificing reliability or user control.

### Personalise beyond message wording

Test more differentiated experiences where repeated segment differences justify them.

### Run longer

Use a larger sample and one principal behavioural outcome over a longer period.

**Evidence behind the product:** 418 questionnaires · 38 interviews · participatory design · focus group · usability testing · one-week field test · four-week field evaluation.

**Measurement boundary:** a manually recorded activity represented product-recorded running, not every run a participant completed.

## IN A GLANCE

# When you urgently need a restroom, uncertainty is the real problem

As a runner in Paris, I found myself planning routes around known public toilets, searching well before the need became urgent, or returning home rather than moving away from a reliable facility.

The opportunity was not simply to put toilets on a map. It was to help someone answer, quickly and independently:

**Where can I go? Is it open? Is it accessible? Can I get there now?**

## From recurring problem to functioning private beta in ~7 weeks

## CORE JOURNEY

Find nearby options → Compare practical information → Inspect one location → Navigate externally

## ROLE

Independent product owner

## SCOPE


Paris · mobile web beta

## EVIDENCE

~10 known testers at verified stage

## STATUS

Early usability learning, not market validation



Map discovery and practical place details in the functioning private beta.

## THREE CHOICES THAT SHAPED THE MVP

## Guest access

Urgent discovery should not require registration or trust in an unfamiliar product.

## External navigation

Use an established map service rather than building an inferior routing layer.

## Scope control

Defer notifications and harder-to-maintain venue types until the basic journey is trustworthy.


Independent product ownership

MVP scoping

User-centred prioritisation

Fast iteration

# AI made it faster to build. It did not decide what was ready to release.



## What I owned

**Problem definition:** use case, target journey and MVP boundary.

**Interface direction:** AI-assisted concept exploration with human responsibility for hierarchy, accessibility and requirements.

**Implementation direction:** Lovable, Gemini and later Codex/Git from specification to functioning beta and technical audit.

**Release judgement:** separating visible completion from verified behaviour.

## The release question changed

The beta looked usable. That was not enough.

**From “Does it look finished?”  
to “What has actually been tested, what  
remains unresolved, and what should  
block wider release?”**

## Verified locally

- Version-controlled migrations replayed successfully.
- Database tests passed for status updates, review aggregation, moderation threshold and event storage.
- Frontend tests and the production build passed.

## Still unresolved

- Protected analytics Edge Function ingestion.
- Some end-to-end role, permission and deletion checks.
- Data provenance and licence review.
- Retention validation.
- Independent security and privacy review.

**Product lesson:** AI can accelerate design and implementation. Product ownership still means defining the problem, deciding what matters, verifying generated systems and refusing to call unresolved work “done”.

# What if the hardest healthcare decision happens before the booking funnel?

Doctolib becomes much easier to use once a person already knows which practitioner to search for, which consultation reason applies and roughly how the French healthcare pathway works.

For someone unfamiliar with that system, the difficult questions often come first:

Which professional do I need?

Do I need a referral?

Does this practitioner accept new patients?

What will this actually cost me?

1

## Uncertain need

Is sports medicine the right entry point?

2

## External sense-making

Google, AI tools or another person.

3

## Search in Doctolib

Generic filters and multiple profiles.

4

## Interpret the system

Referral, sector, reimbursement and new-patient status.

5

## Book with residual uncertainty

Select a reason and hope the choice is appropriate.

Current journey: the decision starts before the booking funnel.

## Product diagnosis

The booking interface solves appointment access better than it solves uncertainty about what to book.

### Evidence boundary

**Evidence:** one structured walkthrough + first-hand experience + eight anecdotal accounts.

**Supports:** hypotheses worth researching.

**Does not support:** prevalence, causality or proof that the proposed solution works.

### What this case demonstrates

Product diagnosis

Behavioural segmentation

Evidence discipline

Responsible design

# Add clarity before adding more guidance

The recommendation is deliberately staged: lower-risk clarity first, then more consequential guidance only after research and validation.

## 1 · ORIENTATION

### New to the French healthcare system?

Optional explanation of médecin traitant, parcours de soins and reimbursement basics.

## 2 · ADAPTIVE SEARCH

### Sports-medicine filters

Clearer new-patient status and structured, factual expertise fields where reliable data exists.

## 3 · DECISION SUPPORT

### Before you book

Separate referral relevance, listed price, reimbursement basis and uncertainty.

## 4 · GUIDED HELP

### Not sure which professional?

Administrative guidance and relevant categories without diagnosing the user.

## DO NOT LAUNCH ALL OF IT AT ONCE

### RELEASE 1

#### Lower risk / broad value

New-patient clarity  
Referral & reimbursement definitions  
First-consultation clarity

### RELEASE 2

#### Data-dependent search

Speciality-adaptive filters  
Structured expertise fields  
Factual profile comparison

### RELEASE 3

#### Higher safety burden

Optional guided selection  
System orientation  
Clinical/legal/accessibility review

### EXPLORE

#### Longer term

Waiting lists  
Insurer integration  
Personalised cost estimation

### How I would know whether it helped

**Measure:** booking completion · backtracking and external search · comprehension · confidence before booking

### Guardrails

No diagnosis · no reimbursement guarantee · no added delay for urgent care · monitor provider fairness and accessibility

## SYNTHESIS

# What I bring to product teams

The value is not the methods by themselves. It is using evidence to make better product decisions.

**Understand the behaviour behind the metric**

Go beyond “where is the drop-off?” to investigate motivation, context, friction and meaningful differences between users.

**Turn insight into product decisions**

Translate evidence into requirements, prioritisation, flows, feature logic and scope choices.

**Carry decisions into delivery**

Work across Design, Engineering and Research, manage trade-offs, and protect the core product mechanism when constraints appear.

**Test whether the decision actually worked**

Separate positive user response from behavioural change, implementation success from product success, and hypothesis from validated evidence.

## EVIDENCE ACROSS THE CASES

**MRunners**

**10/10 frustrated by a challenge** → **redesign**. Unreliable Garmin integration → manual entry. Notification recovery → 27/27 planned notifications verified across six account/device configurations.

**Royal Rest Finder**

**Urgent-use problem** → **guest-first MVP**. Functioning beta → technical verification and explicit release blockers rather than a premature production-readiness claim.

**Doctolib**

**Booking friction** → **diagnosis of pre-funnel uncertainty**. Concept → staged validation before higher-risk guidance.

**I bring behavioural science and research depth, but I use them in service of Product decisions: understanding the problem, choosing what to build, managing trade-offs and testing whether the decision actually worked.**