

AI as a Thinking Partner: Decision Navigator Case Study

Meghna Bhagwat | Product Manager | [LinkedIn](#) | Live Tool: [Click here](#)

PROBLEM STATEMENT

Context: Job hunting is hard, and not just because of the rejection. At every turn, I find myself making high-stakes decisions under real time and financial pressure, with very little information and no clearly "right" answer. Some days it feels like I'm stuck between two impossible choices, other days I'm drowning in too many options. Either way, the pressure doesn't just make decisions harder, it makes thinking harder.

Where I feel this most:

- Do I keep applying, or step back and invest time in building skills that make me a stronger candidate?
- How do I pursue an opportunity in front of me today without selling myself short on long-term fit?

Tracking applications and certifications in Excel is easy, but a spreadsheet can't help you think through uncertainty. It tells you what you've done, not what to do next.

Problem: How might AI help structure thinking and reduce decision fatigue for ambiguous, high-stakes personal decisions without ever making the final decision itself?

USER

Me, someone in the thick of a job search, making frequent, mentally taxing decisions with real consequences and no clear playbook.

WHY AI

What drew me to AI was how it thinks. It reasons through problems making it a natural fit for ambiguous situations. I needed a thinking partner, not a decision maker.

PRODUCT PRINCIPLES

- Lead with clarity, not conclusions.
- Bring structure to the thinking, not advice.
- Favor reversible paths over premature commitment.
- Build calibrated confidence, not false certainty.
- The human always makes the final call.

SOLUTION OVERVIEW: DECISION NAVIGATOR

What it does:

- Reframes the decision so I'm solving the right problem.
- Surfaces hidden assumptions and constraints I might be overlooking.
- Compares options using structured, logical frameworks.
- Defines risks and criteria worth reviewing before I commit.
- Recommends portfolio-style strategies where relevant.

What it doesn't do:

- Make the final decision.
- Rush me toward an answer.
- Make predictions about outcomes.

DESIGN

Phase 1 — Clarifying questions: Before any analysis, the Navigator reads the decision as I describe it and generates targeted follow up questions specific to that decision, up to 4. It stops asking as soon as it has enough to move forward, and if it still doesn't after 4, it says so directly rather than forcing an analysis from thin input.

Phase 2 — Structured output:

- **Decision Reframe** — Am I solving the right problem?
- **Key Assumptions** — What am I taking for granted?
- **Constraints** — What's genuinely limiting my options?
- **Reversible vs. Irreversible Factors** — What can be undone, and what can't?
- **Options Analysis** — A clear comparison of the paths available
- **Recommended Strategy** — A reasoned direction, not a decision
- **2–4 Week Action Plan** — Concrete next steps to keep things moving
- **Review Criteria** — How will I know if it's working?
- **Confidence Level** — An honest read on how certain the analysis really is

PROMPT AS A PRODUCT SPEC

Rather than just writing instructions, I approached the system prompt like a PRD. Each rule targeting a specific failure mode I wanted to avoid such as jumping to advice too early, projecting false certainty, or letting my emotional bias cloud the reasoning.

RISKS & MITIGATIONS

Risks:

- Jumping to conclusions before the full picture is clear.
- Reinforcing my existing beliefs rather than challenging them.
- Sounding more confident than the analysis warrants.
- Reading too much into weak or incomplete signals.

Mitigations:

- Clarifying questions are mandatory before any analysis begins.
- Options are compared side by side through a structured framework, rather than reasoned down a single path, to avoid reinforcing one existing belief.
- Every recommendation carries an explicit confidence label.
- Recommendations are time-boxed to avoid locking in too early.

EVALUATION FRAMEWORK

Since there's no "correct" answer to measure against, I evaluated on five questions instead of accuracy:

- Did it actually help me think more clearly?
- Was the reasoning easy to follow?
- Did it feel honest and reliable?
- Did I feel less overwhelmed after using it?
- Did it know when to step back and let me decide?

I tested the live version against these questions across 10 real decisions, 4 of my own and 6 run by close friends, 7 people total. Testers were not PMs, this was personal decision testing, not a professional evaluation.

Results:

- Helped think more clearly: yes, consistently.
- Reasoning easy to follow: yes, clear and direct.
- Felt honest and reliable: not yet clear, no strong read either way.
- Felt less overwhelmed: partially, it didn't remove the weight of the decision, but helped clear some of the fog.
- Knew when to step back: yes, most confidently. It never jumped to a conclusion across all 10 runs, always asked clarifying questions first, capped at 4, and consistently returned options and a recommended strategy, not a decision. On a deliberately vague input, it correctly recognized insufficient information after 4 questions and asked to start over. On a well-explained input, it recognized sufficiency early and moved to analysis without asking all 4 just to hit a quota.

On average, each analysis surfaced about 5 assumptions and 4 constraints per decision.

LEARNINGS

What worked:

- Reframing the decision upfront cuts through the anxiety and gets me focused on the right problem.
- Structured review cycles keep me moving forward instead of going in circles.

What didn't:

- Vague inputs lead to vague outputs. This directly shaped the live tool's redesign, replacing a single open input with a capped, dynamic clarifying-question flow built specifically to surface the missing context before analysis begins.
- Longer prompts muddy the thinking rather than sharpen it.