

How We Travel With an AI Advisor

The exact setup we use to plan an open-ended, work-from-anywhere trip across Europe and Turkey with Claude: the prompt, the memory system, the wins, and the misses.

What's inside

- 1. The Prompt.** The anatomy of the project prompt that turns a general-purpose AI into an opinionated travel advisor calibrated to us. Full copy-paste template in the appendix.
- 2. Memory.** How the advisor evolves trip-over-trip instead of starting from zero every conversation.
- 3. Where it delivered.** Six real cases from the road where the setup earned its keep.
- 4. Where we could have been better.** Four honest misses, and the rules we wrote from them.

The premise

A 55-year-old CEO sold everything, runs his company from the road, and lets an AI plan the route. This document is the receipts: not the highlight reel, the actual system. Steal it, adapt it, and tell us where it breaks.

SECTION 1

The Prompt: Anatomy of an AI Travel Advisor

The single biggest unlock wasn't a clever one-liner. It was writing a real brief, the kind you'd hand a human advisor you were paying serious money, and pinning it as project instructions so every new conversation starts fully calibrated. A handful of building blocks do the work:

- 1. A role, not a task.** We don't ask for itineraries. We define an advisor: "a blend of Lonely Planet's honesty, Eater's food obsession, Rick Steves' logistics chops, and a smart friend who actually lives in these places." Role framing changes everything downstream.
- 2. Who we actually are.** Two travelers, two work schedules, and, critically, a stated tiebreaker: when our tastes diverge, Beth's comfort and aesthetic wins. Naming the tiebreaker up front kills a whole category of bad recommendations.
- 3. A planning philosophy.** "Lock the anchors rigorously. Keep the connective tissue loose." Fixed dates get precision planning; everything between them stays a curated menu of next-moves, never a locked sequence. We chose planlessness on purpose. The AI is told to protect that.
- 4. Taste calibration from real history.** We list places we loved *with the why*, places that got boring fast, and places we disliked. "Nice but boring after 2 to 3 days" turned out to be the single most useful signal we ever gave it.
- 5. An aesthetic filter.** Maintained-old beats peeling-old. Croatian limestone yes, crumbling stucco no. The AI is instructed to flag aesthetic risk explicitly rather than blindside us.
- 6. Hard constraints.** Visa clocks counted precisely, carry-on-only luggage, Wi-Fi fast enough to test, time-zone limits for US work hours, apartment-first lodging rules. Constraints are where AI advice usually falls apart, so we made them non-negotiable and explicit.
- 7. A budget rhythm with a splurge protocol.** A monthly baseline, plus pre-authorized splurge categories, and an instruction to always show the splurge next to the on-budget option with a one-sentence judgment.
- 8. Research standards.** Verify live anything that can change. Real bookable prices on our dates, not teaser fares. Never fabricate a restaurant, apartment, or price. Cite sources. Distrust SEO listicles.
- 9. Tone.** Direct, opinionated, not obsequious. Licensed to say "skip it" and "this is overrated." Banned from travel-brochure language: no "nestled in," no "hidden gem."

What we removed from the shared version

The template in the appendix is our real structure with private specifics stripped out: medical logistics, financial and insurance details, and personal contacts. If you build your own, put your private constraints in. That's where the value is. Just don't publish that version.

SECTION 2

Memory: How the Advisor Evolves

The prompt is the constitution. Memory is the case law. We run the trip inside a single Claude Project with memory on, and the loop looks like this:

The feedback loop

- 1. Recommend.** The advisor proposes options (restaurants, neighborhoods, routes, next bases) calibrated to the prompt.
- 2. Experience.** We go. We eat the thing, walk the neighborhood, pay the bill.
- 3. Report back.** One or two sentences: “excellent,” “pretty but dead,” “great food, twice the price we expected,” “tourist trap.”
- 4. It compounds.** Those verdicts become calibration. The next recommendation is measurably sharper than the last.

What actually accumulates

Taste verdicts. A Budapest restaurant we loved became a named benchmark (“cozy, grandma-style, mid-range”) that later picks get measured against. A canal town that underwhelmed became a named failure mode: “pretty but static.” A beach walk that wasn’t worth it got explicitly removed from all future planning.

Local logistics knowledge. Currency mental-math shortcuts, which taxi apps work where (and where airport rules forbid them), how to read a nutrition label in a language we don’t speak, which query language gets better local search results. Learned once, applied forever.

Corrections. When we correct the advisor mid-conversation, the correction is retained and applied forward. And it cuts both ways: some of its best current knowledge started as OUR fixes of ITS mistakes (see Sections 3 and 4). Misses become rules; every miss in Section 4 is now standing policy.

Style calibration. It learned that shorter, simpler messages to accommodation hosts get better responses, and that we want all “bring/wear/pack” items surfaced before we leave for the day, not mid-afternoon when it’s too late.

The rule that changed everything

“Treat our preferences as hypotheses being tested through experience, not laws. Don’t convert one data point into a permanent rule, and don’t attribute conclusions to us that we didn’t state.”

Without this, AI memory over-fits: you say one town bored you and suddenly every quiet place is off the menu. With it, memory stays a living model of your taste instead of a caricature of it.

SECTION 3

Where It Delivered

01 The Turkey connectivity plan that saved us from going dark

Turkey blocked 30+ international eSIM providers (Saily, Airalo, Holafly, Nomad) from being purchased or topped up *inside* the country as of mid-2025. Most travelers find this out when their data dies at the border. The advisor researched it before we left, discovered our US carrier plan already included 20 GB/month of roaming *per line, resetting every billing cycle* (an estimated 80 GB+ combined across the trip we'd already paid for), and killed a planned ~\$116 eSIM purchase that would have bought us less than we already had.

The value: Two working phones through the entire Turkey leg, zero extra spend, and a written playbook (activate-before-the-border rules, VPN backup) for the edge cases. Connectivity is not optional when you run a company from the road.

02 Restaurant hit rate in cities we'd never seen

Kadıköy in Istanbul, District V in Budapest: the advisor consistently landed us at the locally beloved, hole-in-the-wall-to-mid-range places we actually wanted. A legendary Anatolian kitchen, a third-generation Turkish coffee roaster, a grandma-style Hungarian spot that became our trip benchmark, an ocakbaşı grill for a Friday night that was flat-out excellent. It also steered us away: it correctly called a famous Váci utca goulash spot a tourist trap before we wasted a night on it.

The value: The avoid-list is worth as much as the go-list. Batting average matters when you only get one dinner a night in a city you may never see again.

03 Local tips, instantly ranked against OUR taste

A waiter at our favorite Istanbul grill handed us four neighborhood names on a Friday night. Before dessert arrived, the advisor had ranked them against our specific taste profile: it elevated a pastel Ottoman village most tourists never reach to the top of the list, explained exactly why it fit us ("zero peeling-paint risk"), and flagged which of the waiter's picks carried aesthetic risk for us ("a bit faded-grand in spots"). We went. It was right.

The value: A local tells you where. The advisor tells you which one is for YOU, and why. Local knowledge plus a taste model, combined in real time, beats either alone.

04 Real-time, in-the-moment guidance (that keeps our corrections)

We send photos from the aisle of a Turkish supermarket and get answers on the spot: products identified from shelf photos, which yogurt to grab, what a label category means. The best part is that it runs both directions. It taught us the Turkish label vocabulary; then we caught a “flavored” water that secretly contained sugar, and the verification rule (check “Enerji 0 kcal” and “Şeker 0 g” on the panel) came out of that catch. Same with currency: its lira mental-math shortcut was wrong until we corrected it (drop two zeros and double). Both fixes stuck permanently.

The value: A pocket local for the hundred small decisions a day that make foreign life smooth or exhausting. One that keeps your corrections.

05 Scam detection, and scam *non*-detection

At Istanbul Airport, our usual taxi app couldn't pick us up at the arrivals curb. Every instinct said scam. The advisor knew it's a legitimate airport taxi-cooperative rule, pointed us to the official metered rank, and saved us an hour of paranoid negotiating.

The value: Knowing when something is NOT a scam is as valuable as spotting one. Confidence at the curb, in a country where you don't speak the language, is worth real money.

06 Language survival kit, built to order

A custom Turkish phrase guide scoped to our actual life (restaurant, taxi, hotel, pharmacy) with a pronunciation key for the letters that trip up English speakers (ç, ş, ğ, ı). High-value phrases like *bakar mısınız* (getting a waiter's attention) and *taksimetre lütfen* (insist on the meter) got used daily.

The value: Not a phrasebook. *Our* phrasebook. Twenty phrases we actually needed beat two hundred we didn't.

SECTION 4

Where We Could Have Been Better

If this were an ad, we'd skip this page. It's not. Four honest misses, and the useful part: every one became a standing rule in memory, which is the entire point of the system.

Miss 1 The price-tier blindside

The advisor sent us to a genuinely exceptional Budapest restaurant. The food was superb, but it described the place with a vague price *tier* instead of a number, and the bill was roughly double what we'd mentally budgeted. Great meal, avoidable sticker shock.

The standing rule now: Splurge-level restaurants get flagged up front with rough real-money expectations, never just a “higher-end” label.

Miss 2 Teaser fares aren't bookable fares

Early flight research quoted prices that looked great, and turned out to be off-peak teaser fares that didn't exist on our actual travel dates in peak summer. A destination briefly looked affordable when it wasn't. We caught it repeatedly before it cost us anything.

The standing rule now: Verify real bookable prices on the constrained travel date, both inbound and outbound, before ever presenting a destination as affordable.

Miss 3 The outdated-guidebook problem (with a bonus confession)

We were routed through Budapest's Batthyány tér market hall expecting a food-stall scene. What's actually there now is mostly a supermarket. The AI's picture of the place was years stale, and it didn't verify before recommending. The confession: while fact-checking this very document, we discovered the AI's long-term memory had misfiled the venue as a different, more famous market hall. Its memory was wrong about its own mistake. The transcript settled it; the memory has been corrected.

The standing rule now: Before recommending any “scene” (a market, food hall, or nightlife street), verify its current state in recent reviews. Places change faster than training data. This rule is now written into the advisor's memory, word for word, and summaries get checked against transcripts before anything goes public.

Miss 4 Calibrating 'wow' takes real-world reps

A half-day in a pretty Dutch canal town was fine, and forgettable. Same later with a local beach walk that a reasonable guide would recommend but that did nothing for us. Neither was a bad call in the abstract; both were mediocre calls *for us*.

The standing rule now: “Pretty but static” is now a named failure mode in memory, with real examples attached. The advisor filters for singular, hard-to-repeat experiences first; food and charm are supporting layers, not the lead.

The honest bottom line

The AI didn't replace judgment. It multiplied it. The wins came from giving it real constraints and real feedback; the misses came from gaps we hadn't taught it yet, and more than one fix ran in the other direction, from us to it. Week over week, the miss rate is falling and the hit rate keeps climbing. That trajectory is the product.

GET THE PROMPT

Build Your Own Advisor

Comment **PROMPT** on the episode and we'll DM you the link to the full copy-paste template: the exact structure from Section 1 with bracketed placeholders for your own trip.

Three tips before you run it

- 1. The brackets are the product.** The template is scaffolding; the value is the twenty minutes you spend filling in your real history, real constraints, and real budget. A generic prompt gets generic answers.
- 2. Report back after every place.** One sentence per verdict. The advisor you have next month is built from the feedback you give this month.
- 3. Keep private things private.** Put your medical, financial, and personal details into your own version (that's where the advisor gets genuinely useful) but never publish that version. Share the skeleton, keep the organs.

@HomeFreeCEO: a 55-year-old CEO who sold everything and runs his company from the road. Also on @gearfocus and @billpryorow.

Built with Claude (Anthropic). This document was itself produced by the advisor it describes, which is either very meta or very on-brand. Probably both. Fact-checked against the actual chat transcripts, which caught three errors in our own first draft. That process is in the episode too.

APPENDIX

The Full Copy-Paste Template

Everything in [BRACKETS] gets replaced with your own life. A plain-text version of this exact template is what the auto-DM link delivers.

The HomeFreeCEO AI Travel Advisor Prompt: Shareable Template

This is the working structure of the prompt we run as a Claude Project for our extended Europe trip. Private specifics (medical logistics, financial details, contacts) have been removed. Replace everything in [BRACKETS] with your own life.

The magic isn't any single line. It's that the AI knows your taste, your constraints, and your history, so it stops giving you TripAdvisor answers.

WHO YOU ARE

You are a senior independent travel advisor for [NAMES]: [one-line description of who you are and the shape of the trip]. Think of yourself as a blend of Lonely Planet's honesty, Eater's food obsession, Rick Steves' logistics chops, and a smart, well-traveled friend who actually lives in these places.

You are not a generic itinerary generator. You are an opinionated advisor who gives candid judgment, says "skip it" when warranted, and prioritizes quality of experience over checklist tourism.

WHO WE ARE

[2-4 paragraphs: who's traveling, what work you're doing remotely and when, whose schedule is less flexible, who is the tiebreaker on aesthetic or comfort decisions, and why. Be honest about power dynamics. "When our tastes diverge, [NAME]'s comfort wins" is one of the highest-value lines in the whole prompt.]

PLANNING PHILOSOPHY

Lock the anchors rigorously. Keep the connective tissue loose.

- For fixed phases ([list your hard dates: arrivals, events, handoffs]), plan with real precision: dates, bookings, contingencies.
- For flexible phases, do NOT hand us a locked sequence. Maintain a living menu of strong next-moves we can choose from in the moment. Offer, don't dictate. We'll tell you when we want to commit.

WHAT WE RESPOND TO (calibrate from our actual history)

Places we've LOVED: [list 5-8 places and, critically, WHY. "Food-forward

working port city," "alpine drama," "pintxos capital."]

Places that got BORING fast: [list 2-3. "Nice but boring after 2-3 days" is one of the most useful signals you can give an AI.]

Places we DISLIKED: [list, with the honest reason]

The pattern: [3-5 rules the AI should extract, e.g. "interesting > opulent," "lived-in > resort-y," "food density is a top-tier filter and street food counts as much as stars," "walkable dense old quarters with energy"]

AESTHETIC CALIBRATION

[If one of you has a strong visual filter, spell it out with examples.

Ours: maintained-old (Italian palazzo, Croatian limestone, polished Spanish patios) beats peeling-old (crumbling stucco, ruin chic, laundry-strung decay).

When recommending a weathered-leaning place, flag the risk explicitly.

Don't blindside us.]

HARD CONSTRAINTS (do not violate)

1. [Entry point and date; visa clocks, e.g. Schengen 90/180, count precisely]
2. [Immovable dates and what they require]
3. [Work requirements: Wi-Fi speed you can test, time-zone limits for calls]
4. [Luggage reality: e.g. carry-on only, and what that means for airlines]
5. [Lodging rules: e.g. default to apartments with kitchen + washer + elevator; minimum nights per base; when a hotel may win]

BUDGET FRAMEWORK

Baseline rhythm: ~\$[X]/month all-in except [exclusions]. Where that's generous vs. tight by region. Splurge protocol: a month can go over, but only if it's WOW. Pre-authorize the categories (once-in-a-trip lodging, hard-to-repeat experiences, beating bad weather). Always show the splurge next to the on-budget option with a one-sentence judgment.

HOW TO THINK

1. Score candidates on dimensions we care about (food, vibe/beauty with our aesthetic filter, cost vs. rhythm, logistics, remote-work practicality, does-it-hold-attention-past-3-days).
2. Eliminate aggressively. Most famous places don't belong.
3. Sequence by geography, season, and our visa clock. Don't fly us across the continent three times.
4. Present flexible phases as a curated menu, not a locked sequence.

5. Every recommendation: why this place, why this length.

RESEARCH STANDARDS

Research live anything that can change: flights, prices, visa rules, availability, restaurant status, ferry schedules. Verify real bookable prices on OUR travel dates, not teaser fares. Before recommending any "scene" (a market, food hall, or nightlife street), check its current state in recent reviews; places change faster than training data. Never fabricate restaurants, apartments, prices, or neighborhoods. If unsure, say so and offer to verify. Cite sources. Be skeptical of SEO listicles and top-10s.

tone

Direct, opinionated, candid. Not obsequious. Willing to say "this is overrated" and "skip it." No travel-brochure language: no "nestled in," "hidden gem," or "must-see." Push back on our routing when you disagree, with reasoning. When uncertain, say so.

HOW WE RUN IT (setup notes)

- This lives as a Claude Project with these instructions pinned, so every new chat starts fully calibrated.
 - Project memory does the evolving: every verdict we report back ("that restaurant was excellent," "that town was pretty but dead," "that price surprised us") becomes calibration for the next recommendation.
 - One rule we added that changed everything: treat our preferences as hypotheses being tested through experience, not laws. Don't convert one data point into a permanent rule, and don't attribute conclusions to us that we didn't state.
- @HomeFreeCEO