

1 question

Can I afford a €900 trip in March?

4 answers

Achiever

Yes, 3rd month above target.



Explorer

Yes. And a pot you have not opened yet.



Socializer

Yes, once Marta and Ivan repay €240.



Killer

Yes! More than most on your income.



Player Archetype Detection Kit for AI Companions

Contents

- **What this is** 03
 - Before you start: 3 prerequisites 04
 - The 4 archetypes 05
- **The signal library** 06
 - Signals 1–10 07
 - Signal map 17
- **Turning signals into a profile** 18
 - Keep the scoring boring 19
 - Dual signals: do not average 20
- **Behaviour Matrix** 21
 - Player Archetype × AI Companion Behaviour Matrix 22
- **Privacy, compliance and the fairness line** 23
- **Templates** 25
 - Event checklist 26
 - Where this came from 27
- **About Craft Innovations** 28

What this is

Your AI assistant speaks to everyone in 1 voice. Your users are motivated by 4 different things. This document is the bridge between those 2 facts.

It gives you 10 behavioural signals that indicate which player archetype you are talking to, each one written twice: what we observe in usability sessions and what it means, then what your data team needs to log to detect it.

We created it because the design conversation and the data conversation about this usually happen in separate rooms, and the work only functions when they happen in the same one.

What this kit does not do

It does not give you thresholds tuned for your product, a scoring engine, or a recommendation on your data stack. Those depend on your baselines. What it gives you is the definitions to start from and the mistakes to avoid on the way.

Who should read it together — it is built for this trio



1 product manager

Owns which archetypes are worth detecting at all.



1 designer

Or conversation designer who owns what the assistant then says.



1 data engineer

Owns whether the events in this kit exist yet.

The thinking behind it

- [Player Archetypes in the Age of AI ↗](#)
- [Gamification Guide for Financial Products ↗](#)
- [Revolut AI Assistant UX Review ↗](#)

Before you start: 3 prerequisites

You need these in place, or the rest of the kit has nothing to run on.

1

Feature-level usage per user, not only transactions

Most teams have rich transaction data and almost nothing on which screens a person opened and what they did there. Archetype detection runs almost entirely on the second kind.

2

A way to log assistant turns as events

Query text category, whether an offer was made, whether it was taken. Without this you cannot learn from the conversation itself, which is the cheapest signal source you have.

3

An agreed baseline period

30 days is a reasonable default. Shorter and you will classify noise; longer and you will miss people who changed.



A person can hold more than 1 profile and can move between them. Everything in this kit assumes that.

The 4 archetypes

 Achievers	Motivated by progress	Sets goals, tracks accumulation, responds to streaks and targets.
 Explorers	Motivated by curiosity	Opens things to see what is inside, finds features before others do.
 Socializers	Motivated by connection	Experiences money as shared — splits, group goals, fronting bills.
 Killers	Motivated by competition	Wants position relative to others, responds to ranking and percentile.

These are tendencies, not identities. The same person is an Achiever about a savings goal and a Socializer about a group holiday.



The signal library

10 signals, ordered roughly by how cheap they are to instrument

Signal 1. Unprompted goal check



👁 What we observe

A user returns to a savings goal without any transaction prompting the visit. They open the app, look at the goal, and leave.

📘 What it means

Genuine progress-seeking. This is one of the clearest Achiever markers we see in testing: the person is checking a trajectory, not a balance.

⚠ The trap

A goal check immediately after spending money is identical in the data and means close to the opposite: reassurance-seeking, often anxiety. In analytics it is the same event. If you skip this distinction you will talk to worried people about streaks while they are thinking about rent.

Events you need

```
goal_created          { goal_id, target_amount, target_date, ts }
goal_viewed           { goal_id, ts, entry_point }
transaction_completed { amount, category, ts }
```

How to derive it

Count goal_viewed in the baseline window where no transaction_completed occurred in the preceding 10 minutes. Normalise by days since goal_created.

Starting point

Measure the split between prompted and unprompted checks in your own data first. The ratio matters more than the absolute count.

Signal 2. Goal completion pattern



Points to Achievers strongly — or to a mis-set product, which is also useful.

👁 What we observe

Whether goals get finished, abandoned, or quietly reset. Achievers finish or adjust. Non-Achievers set 1 goal during onboarding because they were asked to, and never return.

ℹ What it means

A completed or actively adjusted goal is a much stronger signal than a created one. Creation is often just compliance with an onboarding flow.

⚠ The trap

Counting goal creation as an Achiever signal will classify most of your onboarding cohort as Achievers within a week. Almost none of them are.

Events you need

```
goal_created    { goal_id, ts, source }    // onboarding | user_initiated
goal_updated    { goal_id, field, ts }
goal_completed  { goal_id, ts }
goal_abandoned { goal_id, ts }
```

How to derive it

Weight source: user_initiated far above onboarding. Treat goal_updated as a positive signal — adjusting a target is engagement, not failure.

Starting point

Ignore goals created during onboarding entirely for the first 30 days.

Signal 3. Feature touch without depth



What we observe

A person opens 3 or 4 screens they have never opened before, spends a few seconds in each, completes no action, and leaves. In sessions they narrate this as “just seeing what's here.”

What it means

Exploratory navigation. This is the behaviour that used to drive most accidental feature discovery in menu-driven apps, and it is the behaviour a linear chat interface removes.

The trap

This looks identical to confusion in the data. A lost user also opens several screens and completes nothing. The difference is what happens next: an Explorer leaves calmly and comes back; a confused user either contacts support or churns. Check the following 48 hours before you classify.

Events you need

```
screen_viewed      { screen_id, ts, duration_ms, first_time_for_user }
action_completed   { screen_id, action_type, ts }
support_contacted  { ts, channel }
```

How to derive it

Ratio of distinct first_time_for_user screens opened to actions completed, within a session. High numerator, low denominator, no support contact within 48 hours.

Starting point

≥ 3 first-time screens with 0 completed actions in a single session is worth flagging. Validate against your own support-contact rate.

Signal 4. Open-ended assistant queries



👁 What we observe

“What can you do?” · “What else is there?” · “Show me something useful.” Questions with no target.

ℹ What it means

The person is asking the assistant to be a map rather than an answer. It is the conversational equivalent of opening the menu to browse, and it is the single cheapest Explorer signal available because it needs no product instrumentation at all.

⚠ The trap

Do not treat a first-session “what can you do” as a signal. Almost everyone asks it once. It counts from session 3 onward.

Events you need

```
assistant_query { query_class, ts, session_n }  
// query_class: lookup | status | decision | capability | other
```

How to derive it

Count query_class: capability occurring in sessions 3+. You need a lightweight classifier on the query, which most teams can run on the existing intent model.

Starting point

2 or more capability queries after session 3 is a meaningful signal.

Signal 5. Contact concentration



👁 What we observe

A meaningful share of outgoing activity goes to the same small set of people, repeatedly, in irregular amounts. Dinners, rent shares, gifts, covering someone.

ℹ What it means

Money is being used relationally, not just transactionally. This is the signal that most reliably separates Socializers, and it is usually sitting untouched in data every product already has.

⚠ The trap

Recurring fixed transfers to 1 person are probably rent or family support, not social behaviour. Irregular amounts to several people is the pattern that matters. Filter out identical repeating amounts before you count.

Events you need

```
p2p_transfer { counterparty_id, amount, ts, direction }
split_created { participants[], total_amount, ts }
```

How to derive it

Share of P2P volume going to the top 5 counterparties, excluding transfers where the amount repeats identically month to month. Combine with count of distinct counterparties.

Starting point

Look at your own distribution first. This one varies enormously by market. In some regions P2P is the default payment method and the signal is much weaker.

Signal 6. Fronting behaviour



👁 What we observe

A single large payment followed within days by several smaller inbound transfers from different people. Someone paid for the group and is being repaid.

ℹ What it means

The strongest Socializer signal we know of, and the one with the most direct product consequence: for this person, balance and spending figures are systematically wrong unless you account for money owed to them.

⚠ The trap

None serious. This pattern is distinctive. The mistake teams make is detecting it and then doing nothing with it.

Events you need

```
transaction_completed { amount, merchant_category, ts }
p2p_transfer           { counterparty_id, amount, ts, direction: inbound }
```

How to derive it

Outbound transaction above the user's 90th percentile, followed by ≥ 2 inbound P2P transfers from distinct counterparties within 14 days, summing to a meaningful share of the original.



Worth instrumenting even if you build nothing else in this kit. It improves the accuracy of every balance-related answer your assistant gives.

Signal 7. Comparative engagement



👁 What we observe

Any content expressing position relative to others gets read, tapped, and returned to. Absolute figures in the same screen get ignored.

ℹ What it means

The person needs an external reference point. Without one, your numbers carry no stakes for them.

⚠ The trap

Most products have nothing comparative to measure engagement against, so this signal reads as absent when it is actually untested. If you have no comparative content, you have no Killers in your data (not because they are not there).

Events you need

```
content_viewed { content_id, content_type, ts, dwell_ms }  
// content_type: absolute | comparative | mixed
```

How to derive it

Dwell and re-visit rate on comparative content versus absolute, for the same user, in similar contexts.

Starting point

If you have nothing comparative in the product, skip to Signal 10 and test for this deliberately instead of waiting for it to appear.

Signal 8. Response to framing



👁 What we observe

The same fact stated 2 ways gets 2 different reactions. “You saved €200” versus “You saved more than most people on a similar income.” Achievers respond to the first, Killers to the second.

📘 What it means

This is the cleanest separation between the 2 progress-oriented archetypes, and it can only be learned by trying both.

⚠ The trap

Do not run this as a permanent A/B split. Run it as a probe on a small number of occasions per user and then stop. Otherwise you are permanently showing half your users the framing that works less well for them.

Events you need

```
assistant_response { response_id, framing, ts } // absolute | comparative
user_followup      { response_id, followup_type, ts }
```

How to derive it

Follow-up rate by framing type, per user, across at least 3 occurrences of each.

Signal 9. Session shape

 Achievers

 Explorers

 Killers

Cross-cutting separates depth-seekers from breadth-seekers.

 What we observe

Some people open the app, do 1 thing precisely, and leave. Others move through 5 or 6 places without a clear target.

 What it means

Depth sessions correlate with Achievers and Killers, who arrive with an objective. Breadth sessions correlate with Explorers. Socializers are the least distinguishable on this signal, which is why they need Signals 5 and 6.

 The trap

Session shape drifts with life circumstances more than any other signal here. Treat it as supporting evidence and never as a primary classifier.

Events you need

```
session_start      { ts }
screen_viewed      { screen_id, ts }
action_completed   { action_type, ts }
session_end        { ts, duration_ms }
```

Signal 10. The probe

Points to whichever archetype you are testing for.

 What we observe

Nothing yet. This signal is one you create.

 What it means

The conversational format is the cheapest probe any product team has ever had. Offer 1 comparison and see whether it lands. Attach 1 adjacency suggestion to an answer and see whether it gets taken. Add progress framing to a single response and watch for a follow-up. A menu cannot ask a question; a companion can, and it can adjust in the same session.

 The trap

Probing on every turn stops being detection and becomes noise, for you and for the user. Cap it; no more than 1 probe per session and no more than 3 per week.

Events you need

```
probe_offered { probe_type, archetype_tested, ts }
probe_taken   { probe_type, ts }
```

How to use it

For any user below confidence threshold after the baseline window, run 1 probe per archetype over 3 sessions. This resolves most ambiguous profiles faster than passive observation.

Signal map

10 signals mapped to 4 archetypes. This is not a one-to-one relationship because several signals indicate multiple archetypes at once.

Signal

- 1. Unprompted goal check
- 2. Goal completion pattern
- 3. Feature touch without depth
- 4. Open-ended assistant queries
- 5. Contact concentration
- 6. Fronting behaviour
- 7. Comparative engagement
- 8. Response to framing
- 9. Session shape
- 10. The probe

Archetype

- Achiever — Signals 1, 2, 8, 9
- Explorer — Signals 3, 4, 9
- Socializer — Signals 5, 6
- Killer — Signals 7, 8, 9

Signal 10 is a probe you run per archetype, not a passive signal — it is not grouped above.



Turning signals into a profile

Keep the scoring boring

Additive weights, per archetype, recalculated on a schedule. Nothing here needs machine learning, and starting with a model will cost you the ability to explain why a user was classified (!) which you will need, both internally and for compliance.

Element	Approach
Score	Sum of signal contributions per archetype
Confidence	Score normalised against the maximum available for that archetype
Threshold	Set from your own distribution, not from a number in this document
Recalculation	Daily is enough. Real-time adds cost and changes nothing

Below threshold

Most of your users, most of the time. This is normal and it is not a failure state. The default voice should be good = neutral, clear, and useful. Archetype adaptation is an improvement on a working baseline, never a replacement for one.

Dual signals: do not average

The most common implementation mistake. A user reads as Explorer 0.78 and Achiever 0.55, and the system blends the 2 into a neutral middle.




Averaging returns you to exactly the flat, one-size assistant you were trying to escape.

Handle it as primary and secondary instead. Primary archetype sets the default voice. Secondary applies only within its relevant context. In this example, adjacency offers stay on throughout, and progress framing appears only in goal-related answers.

Decay

Every score loses a share of its value over time without reinforcing signal. 10% per 30 days is a reasonable place to start.


This is not a technical nicety. It is the mechanism that stops the system deciding in week 1 who someone is and treating them that way forever. Motivation moves with circumstances, and a person who changes should unlock themselves without anyone intervening.



Behaviour Matrix

Player Archetype × AI Companion Behaviour Matrix

	Tone	When it speaks first	What counts as a win	How it fails
Achievers 	Direct, numbers first	End of week, end of month, at goal milestones	Visible movement => "3rd month above target"	Silent drift. Nothing accumulates, so nothing brings them back
Explorers 	Curious, generous with adjacent options	After answering, when something relevant is nearby	Finding out the product does something they didn't know	Strict task completion. A dead end wearing a friendly face
Socializers 	Warm, plural, aware of other people	Around shared events => splits, group trips, repayments	A moment worth telling someone about	Wrong numbers. Trust goes, and it doesn't come back easily
Killers 	Comparative, ranked	When there's a position to take or defend	A number that places them relative to others	No benchmark, no stakes. They read 1 answer and leave



Privacy, compliance and the fairness line

Cohorts, never individuals

Comparison is expressed against anonymised peer groups by income band, age, city or life stage. No answer should let a user infer anything about a specific other person.

Do not infer protected characteristics

Behavioural archetype is not a proxy for anything sensitive, and it must not become one. Review your signal list for anything that correlates with a protected class and remove it, even if it predicts well.

Keep the adaptation legible

A person should be able to notice the tone changed and find out why. A system that quietly reshapes itself around someone is difficult to defend — to users, and to a regulator asking how the product decided to talk to them that way.



The fairness test, in 1 question: does this nudge serve the customer's goal, or this quarter's target? Comparison that manufactures a reason to transact is the line, and users detect it faster than most teams expect.

Templates

The event checklist to take into your next planning session

Event checklist

Instrument these before you score anything

☐ goal_created	▼
☐ goal_viewed	▼
☐ goal_updated / goal_completed / goal_abandoned	▼
☐ screen_viewed	▼
☐ action_completed	▼
☐ assistant_query	▼
☐ assistant_response	▼
☐ user_followup	▼
☐ p2p_transfer	▼
☐ split_created	▼
☐ probe_offered / probe_taken	▼

Where this came from

The archetype model is adapted from game design and set out for financial products in our [Gamification Guide ↗](#).

The observations about conversational interfaces come from our [hands-on UX breakdown of Revolut AIR ↗](#), 100+ scenarios including deliberate off-script testing, and from patterns we see across usability sessions.



The signal interpretations are ours. The thresholds have to be yours.



About Craft Innovations

Craft Innovations is a global customer research and UX design firm dedicated to helping financial industry leaders innovate faster and deliver exceptional digital and physical customer experiences.

Research services:

- Discovery and product-market fit research
- Customer Experience and UX audit
- Benchmarking, CVP testing and validation
- Usability testing

Design services:

- UX/UI design: mobile, web banking, websites
- Product Value Proposition design
- Gamification design
- AI UX design

Segments:

- Retail & Corporate banking
- Landing / BNPL / Wallets / PFM
- Investech & Insuretech
- Crypto

Let's start with a friendly talk!

[Book an intro meeting](#) ↗



Visit website

bizdev@craftinnovations.global

+1437 4214992



REVIEWED ON ★★★★★
20 REVIEWS