

USER GUIDE

Facilitation (Pronoia)

Delphi ecosystem — user guides

English

Version 1.0 · 2026-09-24

Contact: Hannu Linturi · hannu.linturi@gmail.com

Contents

FACILITATION (PRONOA)

The iterative process: rounds, phases and convergence	5
Why Delphi iterates	5
The Pronoia process at a glance	5
Round 0 — mapping and motivating (AI-assisted intake)	6
The phases inside a round	6
Controlled feedback: the visibility ladder	7
Policy presets — choosing an iteration stance	8
Revision and movement	9
The panelist Home — real-time data that fuels iteration	9
Tracking the panel — process metrics	10
Iterating across rounds — when to continue, when to close	10
The facilitator's iteration playbook	11
Anonymity: identity modes and group reporting	12
Principle	12
Three identity modes (per-panel setting)	12
Cross-cutting rules (always on, in every mode)	13
What is available in each mode	14
Group comparison — current behaviour	14
In practice: role play and the one-shot poll	15
Data model (implemented) — for the developer-partner	15
How Delphi Pronoia's two interfaces are built	16
1. One study, one truth	16
2. The panelist interface	17
3. The facilitator interface	19
4. The same round, from both sides	24
5. What the interfaces deliberately do not do	25
The panelist's process, round by round	26
1. What you are joining	26
2. Before anything opens: the invitation and the consent screen	27
3. Round 0: the orientation conversation	27
4. The frame that never changes	28
5. Your turn to answer	28
6. Answered, and the round is still open	29
7. Dialogue: the round where the panel talks to itself	30
8. Between rounds — the richest screen in the study	31
9. When the study ends	31
10. More services: the index	32
11. Two ways the same journey is named	33
12. If the study is real-time	33
13. Your answers, your data	34

14. For the facilitator: what is worth saying out loud	34
The facilitator's process, from planning to handover	36
0. The map you are working on.	36
1. Planning: a completeness list, not a wizard	37
2. Round 0: orientation, run as a sequence	39
3. Collection: the round-management card	40
4. The response phase: your work is monitoring — and preparing	41
5. The dialogue phase: monitoring and moderation	42
6. The revision phase	43
7. Glancing at results without leaving your work.	43
8. Between rounds: your richest screen too	44
9. Reporting and handover	45
10. Exceptions, and what they cost.	46
11. What the panel sees while you work	47
12. Five things that cannot be undone	47
13. What is folded, and when	48
14. The line that says where on the arc you are	50
15. Retirement is the Barometer's way, not a general one	51
Role play: a role in a box of the panel matrix.	52
What role play is for	52
Division of work: the facilitator gives the structure, the panelist chooses the box	52
The facilitator's steps	53
What the panelist sees.	53
AI roles as panel supplements.	54
Reading the results.	54
Technical description	54
One-shot poll at a seminar: QR code and anonymous answers.	56
What the one-shot poll is for.	56
Before the session	56
During the session	57
Closing the poll	57
Results	58
Limits	58
Technical description	58
Taking part in the open Delphi study	60
No account needed	60
A Delphi study is not a survey	60
Anonymity is structural, not a promise.	61
What else is on the page	61
If something doesn't work	61
Short answers	62
The facilitator trial in five stages	63
1. Open your own example study	63

2. See where the panel landed 63

3. Read the dialogue..... 63

4. Change something 64

5. Build a study of your own 64

The limits of the trial 64

Tell us what didn't work 65

The iterative process: rounds, phases and convergence

Why Delphi iterates

Delphi is not a one-shot survey. Its defining mechanism is **iteration with controlled feedback**: experts answer independently, then receive an anonymized picture of the whole panel's answers and reasoning, and reconsider in that light. Repeating this loop is what separates Delphi from a poll.

The iteration does specific work: it lets reasoning surface and be exchanged rather than just tallied; it gives people a structured chance to revise in response to arguments they had not considered; and across rounds it either moves the panel toward consensus or reveals a **stable, well-understood dissensus** — which is an equally valuable result.

The core sequence is: **independent judgement → feedback → reconsideration**. Anonymity and the staged release of feedback exist to protect that sequence from social pressure — bandwagon effects, dominance and anchoring — so that movement between rounds reflects better reasoning, not the loudest voice.

The Pronoia process at a glance

A study advances through a small state machine that the facilitator drives one step at a time:

Draft → Round 0 → Round 1 → Round 2... → Close → Analysis

Round 0 is the orientation, and it belongs to every variation; every numbered round runs the same internal sequence of phases; and the loop repeats until the facilitator closes the study and hands a clean dataset to analysis. The variation's own character begins when Round 1 opens: Real-Time Delphi runs one round in which responding and dialogue are open at the same time, and Barometer Delphi rates the same theses again from one wave to the next. Classic Delphi and Argument Delphi run the same process; they differ in the aim (consensus / dissensus) and in the indicators, not in the phases.

Round 0 — mapping and motivating (AI-assisted intake)

Round 0 is more than a setup screen. In Pronoia it is run as an **AI-assisted chat dialogue** that does two jobs at once:

- **Motivating.** The dialogue meets people where they are, explains what the study asks of them and why their view matters — so they arrive at Round 1 invested rather than cold.
- **Mapping.** It elicits how each participant frames the phenomenon — the issues, positions and tensions they see — which scopes the argument space and feeds directly into thesis design and the panel matrix.

Selective / key-informant variant. The dialogue does not have to be run with the whole panel. The facilitator can run it selectively with a handful of **key informants** — people who know the terrain — use their input to map the phenomenon and seed the theses, and then invite the wider panel straight into Round 1.

Whichever way it is run, Round 0 ends by fixing the two things that define the argument space: the **panel** (built as the orientation-appropriate matrix, cleared through its quality gates) and the **theses** (at least three, in the structure the orientation calls for). While in this setup state the facilitator can still edit and, where used, regenerate questions; once Round 1 opens, the instrument is **locked** so everyone answers the same questions.

The phases inside a round

Every numbered round runs the same four phases. Separating them is deliberate: it keeps "form your own view" apart from "engage with others", which is the heart of the Delphi method.

Phase	What the panelist does	Sees others?	Interaction	Why
Response (1a)	Answers each thesis independently	Not before submitting	Edit own answer	Capture genuine, un-anchored judgement
Dialogue (1b)	Reads the panel, comments and replies	Yes	Comment + reply, graph visible	Exchange reasoning, not just numbers
Revision	Updates positions in light of the dialogue	Yes	Edit own answer (tracked as a revision)	Let movement happen — and measure it
Closed	Round is locked	Yes (read-only)	None	Freeze a comparable snapshot before the next round

The response and dialogue phases are also called **1a** and **1b** in the facilitator controls — sub-rounds of the same round.

A phase adds a capability, it never removes one. Opening the dialogue phase does not close answering: a panelist who did not get to answer can still answer during the dialogue, and may take part with the full repertoire for the whole round. The idea of Delphi is that a panelist can change their answer when someone else's reasoning gives a good reason — and that requires answering to stay open.

Controlled feedback: the visibility ladder

What a panelist can see is set **per phase**, and it is the main lever Pronoia gives the facilitator over anchoring. Three settings climb from fully protected to fully transparent:

Visibility	Meaning	Effect on anchoring
Blind	Others' answers stay hidden for the whole phase	Maximum protection; use with deliberation
After-submit	No anchoring before you commit, then the panel and the aggregate graph open up	Protected where it matters, feedback where it helps (the eDelphi norm)
Open	The panel is visible throughout the phase	No protection by design; full transparency

Alongside visibility, each phase independently controls whether commenting, replies and the result graph are available — which is how the response phase stays quiet and focused while the dialogue phase opens up the argument.

Policy presets — choosing an iteration stance

Rather than set every phase by hand, the facilitator picks a preset that encodes a methodological stance; individual phases can still be overridden per round.

Preset	Response phase	Dialogue / revision	When to use
Classic	After-submit, no graph	Open	Staged Delphi with anchoring protection — the default
Real-time (RTD)	Open	Open	Continuous, always-open eDelphi; the round never closes for staged feedback
Open	Open, graph visible	Open	Full transparency from the first answer — workshops, teaching, low-stakes panels
Barometer	Open, recurring	Open	The same theses rated again from one wave to the next; tracks how positions move over time rather than converging in one study

Preset and variation are different things. The **Classic** preset is the shared stance of Classic Delphi and Argument Delphi — the same process, a different aim. Real-Time Delphi is a different cadence rather than a different method: instead of discrete rounds, the panel answers and sees feedback continuously for the length of one round, and the facilitator closes when the picture is stable. Barometer Delphi is a monitoring stance: the same instrument is measured again wave by wave, a

thesis may be retired between waves and replaced by a new one, and change over time is itself a finding.

Revision and movement

The point of feedback is to make **informed movement** possible. When a panelist changes an answer during the dialogue or revision phases, Pronoia records it as a revision rather than overwriting silently — so the facilitator can see *how much* and *where* the panel moved between rounds.

That movement is read against the **panel pulse** — the consensus / dissensus picture per thesis. Convergence shows where the panel is settling; persistent dissensus shows where a real, structured disagreement lives. Both are signals for the next decision.

Crucially, the learning movement is **not uniform** across the panel. Different groups — expertise areas, stakeholder positions or observation distances — typically behave differently: some converge quickly, others hold their ground, others shift only late. Reading movement **by group** is central to describing how the panel actually learns — and it is also where the most effective nudging happens: simply surfacing that one group has moved while another has not is itself a prompt for the lagging group to reconsider. Pronoia's group comparison along the panel-matrix axes makes these differential movements visible.

The panelist Home — real-time data that fuels iteration

Iteration only works if panelists come back and engage. The panelist **Home view** is designed to make that happen: it turns the panel's live data into something informative and absorbing, so that seeing the panel move makes the panelist want to move too.

The view adapts to where the panelist is. On the first visit, before they have answered, it stays deliberately minimal — a short "About this study" and a direct route into the questions — so nothing stands between them and their first independent answer. Once they have answered, the layout puts the whole panel first:

See the pulse → Notice movement & tension → Revise → See your effect & learn

Why real-time data fuels iteration. Seeing live consensus form, or a tension open up around one's own position, is what moves a panelist to reconsider and revise. Each revision then changes the picture others see, which prompts further movement — and through that loop the panel, and each panelist, learns.

Tracking the panel — process metrics

Throughout the whole process Pronoia follows the panel with a small set of activity metrics. Together they show, at a glance, whether the panel is alive and engaging:

Metric	What it counts	What it signals
Panelists	Members in the panel	The reach of the study
Respondents	Panelists who have answered	Response rate and coverage
Comments	Arguments and comments written	Depth of reasoning offered
Replies	Comments on comments	Genuine dialogue, not parallel monologue
Reactions	Lightweight responses to others	Low-effort engagement and attention
Revisions	Changes to answers and comments	The learning movement itself
Visits	Visits to the site	Sustained attention and return rate

Revisions and replies are the most telling: they are the traces of panelists actually reconsidering in light of one another — the iteration working as intended.

Iterating across rounds — when to continue, when to close

After a round closes, the facilitator makes one decision: run another round, or stop.

- **Run another round** when positions are still moving, when dialogue raised new considerations that deserve a fresh answer, or when a thin panel has just been supplemented.
- **Close** when the panel has converged, or when dissensus is **stable and well-understood** — further rounds would add cost without insight.

Between rounds the facilitator can also act on panel health: if diversity was too thin, supplement the panel (or add perspective theses) before the next response phase, so the new round genuinely widens coverage rather than repeating the same voices.

Closing hands off to analysis. On close, the collected and validated material becomes the input to the analysis ecosystem (the DAE pipeline and its methods). Iteration in Pronoia ends exactly where structured analysis begins.

The facilitator's iteration playbook

1. **Set the stance first.** Choose a preset (classic / RTD / open) that matches the stakes, and adjust individual phases only if you have a reason.
2. **Open with the intake dialogue.** Run the AI-assisted Round 0 chat — with the full panel or a few key informants — to motivate participants and map the phenomenon before the theses are locked.
3. **Open Round 0, lock the instrument.** Finalise panel and theses, then open Round 1 so everyone answers the same questions.
4. **Advance one phase at a time.** Move response → dialogue → revision deliberately; let the response phase do its un-anchored work before opening the panel.
5. **Keep the panel engaged.** Use reminders and messages, and point panelists to their live Home view — iteration only works if people return.
6. **Read the pulse before deciding.** Look at convergence and dissensus per thesis, and how much positions moved, before choosing to iterate or close.
7. **Tend panel health between rounds.** Supplement diversity where it was thin so the next round adds coverage.
8. **Close cleanly and hand off.** Stop when consensus or stable dissensus is reached, lock the study, and pass the validated dataset to analysis.

Source: Pronoia — *The Iterative Process* (Metodix, June 2026).

Anonymity: identity modes and group reporting

Principle

The core of the Delphi method is to reduce social-conformity pressure and dominance through **(semi-)anonymity**. In Pronoia, panelists are **strictly anonymous to one another yet identifiable through a pseudonym**, so that in dialogue a participant recognises the same conversation partner as the same person. Anonymity protects honest judgement; it must not break accidentally through small groups, AI participants, or comparisons.

Group-level reporting runs along the **panel matrix** — the expertise (primary) and stakeholder (secondary) axes — rather than a single legacy grouping. The same **k≥3 protection** applies to every axis and to the panel-pulse group dots.

Three identity modes (per-panel setting)

A. Pseudonymous (default)

A stable **word handle** per panelist (e.g. Aurora, Vega) — memorable and distinctive, giving no hint of identity or status, stable for the whole study. Real identity sits only behind the invite link (invite token). Full functionality: Home, dialogue, revisions, recognition, arena.

B. Role play (optional layer on top of A)

The panelist **writes a role** or perspective they answer from (e.g. *Municipal finance director in 2035*) and **places it in a box of the panel matrix**. Off by default. The division of work is clear: **the facilitator owns the structure** – the axes and categories of the panel matrix – and **the panelist chooses the box** their role fits best.

- **The role is a name, the box is a group.** Group comparison and analysis use the box the panelist chose, not the free role text. Freely written roles therefore do not break the panel into groups of one, and k≥3 protection stays in force.
- **Others see the role** next to the panelist's reasoning and comments, after the handle, marked *self-selected*. It is a *claimed* perspective, not verified expertise.

- **AI role.** The facilitator can commission a role for an AI panelist in the same matrix. It is shown marked *AI role*, and the AI answers from the role's perspective. A role a human chose and a role an AI plays are never confused.
- When role play is switched off, roles are shown to no one; the saved role and box are kept.

C. Fully anonymous one-shot poll

No identity, a single response – a seminar or mobile poll the facilitator shares as a **QR code** or link. No sign-in, own page, arena, revisions, recognition or persistent handle. Voters answer the theses and may write a short reason; the writer is never named. Each submission is stored as an anonymous participant who does not appear on the panel roster or in group comparison, and the answers go into the round's results and distribution like any other answers.

The duplicate-vote guard is per browser: the same browser answers once. It stops accidental repeats but not deliberate ones – another browser or a private window can answer again. The one-shot poll therefore suits taking the temperature of a room, not a vote whose outcome is binding.

Cross-cutting rules (always on, in every mode)

AI is always disclosed. Every AI panelist is marked with "**AI**" plus a short role description, 2–3 words (e.g. *AI · Economist*, *AI · Labour Researcher*). Shown everywhere an AI panelist appears: dialogue arguments and comments, panel lists, pulse group labels, reports. Never hidden, never presented as human.

k-anonymity (k = 3) across all group-level views. No sub-group with fewer than three people is reported at group level — comparisons, sentiment and panel-pulse group dots included — and the rule applies **independently on each matrix axis**. Complementary protection: when a group falls below the threshold it is folded (into an "Other" group) until every cell has at least three; if "Other" is still below three, the breakdown is not shown at all. Panelists with no value on an axis are excluded from that axis — they never form a residual bucket. The same threshold applies to cross-tabulations (group × thesis): a numeric cell is hidden when $n < 3$. On small panels the facilitator is steered to a 2–3 group split; anonymity wins. In the extreme, the pulse is shown at whole-panel level only.

Visibility & the after-submit reveal. Per-round phase policy controls when a panelist sees others. The default **classic** preset uses **after-submit** for the response phase: no anchoring before you commit your own answer, but once submitted the panelist sees others' anonymized responses and the aggregate graph. Strict **blind** remains available in the facilitator policy editor. All revealed material stays pseudonymous and $k \geq 3$ -protected; the reveal changes *timing*, not the identity rules.

What is available in each mode

Feature	A Pseudonymous	B Role-play	C Fully anonymous
Persistent handle	yes	yes	no
Recognisable in dialogue	yes	yes	no (nameless)
Home / arena / recognition	yes	yes	no
Revisions	yes	yes	no (one-shot)
Commenting	yes	yes	yes, anonymous (moderated)
Self-selected role in a box	no	yes (marked)	no
AI labelling (if any AI)	always	always	always
k≥3 group reporting (both axes)	always	always	always
Facilitator sees person	yes*	yes*	no (no identity)

*Unless blind-facilitator mode is enabled.

By default the facilitator sees real identities (for reminders and quality control); a per-panel **blind-facilitator** mode hides them, so the facilitator also sees only pseudonyms (the invite mapping still exists in the database).

Group comparison — current behaviour

- **Two axes, stacked.** Both the panelist Home and the facilitator Community view show the expertise (primary) comparison and, beneath it, the stakeholder (secondary) comparison — each with its own roster, numeric matrix and sentiment.
- **Numeric comparison where the question type allows.** Group means and their deviation from the panel mean (plus a divergence figure) are computed for scale questions only; other question types do not produce a numeric matrix.
- **Comment-based sentiment for every question type.** Each group's tone (positive / cautious / critical / neutral, with a net figure) is derived from arguments and comments across all theses, so open, ranking and multiple-choice questions are covered too.
- **Pulse group dots** follow the primary (expertise) axis, matching the group-comparison default.

- All of the above remain under **k≥3 with complementary suppression**; no "Out" / "Other" buckets appear for a fully classified matrix panel.

In practice: role play and the one-shot poll

Each has its own working guide:

- Role play: a role in a box of the panel matrix: when role play suits a study, the facilitator's steps, a template for the panel, AI roles, and reading results by box.
- One-shot poll at a seminar: QR code and anonymous answers: preparation, the QR code and presentation view, what to tell the room, disruptions, results and limits.

Data model (implemented) — for the developer-partner

- **Study:** `facilitator_blind`, `role_play_enabled`, `handle_theme`, one-shot fields (`oneshot_comment_window`, `oneshot_moderated`), `policy_preset` + per-round `phase_policy` (`visibility = blind | after_submit | open`), `panel_matrix` (JSON: primary/secondary axis names and categories), `matrix_type` (`expert | observer`), `language`.
 - **Panelist:** `display_handle` (themed, panel-unique, stable), `self_selected_role`, `participant_type` (`human | ai_agent`), `matrix_primary` / `matrix_secondary`, `description`, `is_decision_maker`, `language`.
 - **Axis grouping** (`_panelist_axis`): `primary` → `matrix_primary` (AI perspective as fallback; in role play the box the panelist chose – the role is a name, not a group); `secondary` → `matrix_secondary`; returns empty when the axis is empty, so the panelist is excluded rather than bucketed.
 - **Shared group reporting** (`grouped_display`): `k = 3` + complementary suppression ("Other") + None-exclusion + small-panel fallback to whole-panel level. Used by `/groups` (Home) and `/community` (facilitator) with both axes stacked, and by the pulse group dots.
 - **AI label** derived from `participant_type = ai_agent` plus the `ai_profile` role, rendered as "AI · {role}", role standardised to 2–3 words.
 - **Blind facilitator:** when `facilitator_blind = true` the UI shows only handles (names/emails hidden), though the invite mapping still exists in the database.
-

Source: Pronoia — Anonymity Design (Metodix, June 2026, status: implemented).

How Delphi Pronoia's two interfaces are built

Pronoia has two faces. A **panelist** sees one; a **facilitator** sees the other. They are not two applications and not two versions of the same screen. They are two questions asked of one running study, and the whole design follows from the fact that the two questions are different.

This guide explains the idea and the structure. The two guides beside it walk through the process — what actually happens, in order, with what on the screen.

1. One study, one truth

Both interfaces read the same two facts:

<code>current_round</code>	which round the study is on (0 = orientation, 1...N = Delphi rounds)
<code>current_phase</code>	<code>response</code> · <code>dialogue</code> · <code>revision</code> · <code>closed</code>

Everything visible is derived from those two. There is a third field, `Study.status`, and it is a **presentation** — a label derived from the round and the phase. Nothing decides anything from it.

This sounds like an implementation detail. It is the reason the interfaces are trustworthy, and it was learned the hard way. Before August 2026 the question "*which round are we on*" was answered by five surfaces reading three different sources, one of which recognised rounds by matching the strings `round1`, `round2`, `round3` — so **in a four-round study, round 4 never highlighted anywhere**. Two surfaces could disagree for fifteen seconds because one polled and the other did not.

The rule that replaced it: *round and phase are the truth; anything else is a label*. When you see the same fact in two places in Pronoia now — the round selector and a banner, say — they are two readings of one computation, not two computations.

2. The panelist interface

2.1 The idea

The panelist's question is "**what is open for me now?**" — not "what is the state of the study". The view answers that question directly and organises itself around the answer.

Five principles govern it:

P1 • The phase organises the view, not the menu. What you land on is what is open for you. The facilitator's switches decide exceptions, not the normal case.

P2 • The panel's *process* is always visible; the panel's *content* only after you have answered. "27 of 40 have answered", "two new comments since your last visit" — these say nothing about *what* the panel thinks, so they cannot anchor your judgement. The distribution, the average and other people's arguments do, so they wait until you have answered. The condition is **per thesis**, not per page: answer this thesis, see this thesis's distribution.

P3 • The frame never changes; the main column changes. Same heading, same menu button in the same place, the form always in the same spot. Not different pages — the same page in different states. A returning panelist needs recognisability more than a per-phase optimum.

P4 • An empty service is not in the menu. Every menu row is a promise that something is there. A "Group comparison" that opens empty is worse than no row. Emptiness is computed from the data, not switched on by hand.

P5 • The description and the notices are not extras. The study description sits above the form on first answering; the privacy and AI notices are always reachable. These are part of what the panelist consents to — not content that has to earn its place by being interesting.

2.2 The structure: six states

The panelist view has six states, defined by the panelist's own question:

	State	When	The panelist's question
S0	Not open yet	invited, nothing opened	"What am I coming into?"
S1	Your turn to answer	response phase, no answer from you	"What am I being asked?"
S2	Answered, round still open	response phase, you have answered	"Did that go through?"
S3	Dialogue turn	dialogue or revision phase	"What did the others think — do I change my mind?"
S4	Between rounds	round closed, next not open	"What came of it? What next?"
S5	Finished	study completed	"Where did this lead?"

S4 is the richest moment of the whole process, and it is easy to waste. It is where a panelist sees what the round produced and forms the view they will bring to the next one.

2.3 What is on the screen

The frame is constant: title, **Panel pulse** (how far the round has got), the menu, and the main column that changes with the state.

- **Panel pulse** shows process, never content: how many have answered, what has changed since your last visit, whether the round is live.
- **More services** is an index, not a drawer of extras — and by P4 it only lists what actually exists for this study.
- **AI panelists are marked wherever they have an effect.** If the panel includes AI participants, the panelist is told so, and AI contributions are marked in the distribution as well as in the text. This is not a footnote in the terms; it is a mark on every surface where an AI answer counts.

2.4 Two ways to name the journey

The same structure can be shown in two vocabularies, chosen per study:

	Story	Plain
Round 0	Threshold · Kastalia	Start · preparatory chat
Rounds	Lesche · panel rounds	Rounds · answering and dialogue
Analysis	Temple · Kassotis	Analysis · processing the results

This is naming, not structure: the states, the rules and the surfaces are identical. A study that wants the Delphic imagery gets it; a municipal panel that would find it strange gets plain words.

3. The facilitator interface

3.1 The idea

It is tempting to mirror the panelist view: the study is in the dialogue phase, so show the facilitator the dialogue phase. That is wrong, and the reason is the whole design.

The facilitator's phase is not the study's phase. The panelist asks "what is open for me"; the facilitator asks "**what is my work right now**", and the two do not coincide. While the response phase is open, the study's phase is **response** — but the facilitator's work is *monitoring*, and *preparing the next round*. A straight mirror would give the facilitator a screen that says "wait".

So the unit is not the phase. It is **(round, phase) → the facilitator's task**, and there are four kinds of task:

Study is...	The facilitator's task	Lands in
Draft	building	Planning
Round 0 open	monitoring + intervening	Collection
Round N, response	monitoring	Collection
Round N, dialogue	monitoring + moderating	Collection
Round N, revision	monitoring	Collection
Between rounds	preparing	Analysis
Finished	reporting	Reporting

Plus one thing that is not a place at all: **intervention** — advance, open, close, handle an exception. It is never its own screen; it is always present where it applies.

3.2 The structure: two levels and a bar

Both levels are at the top of the screen, in one navigation that stays in place while you scroll. A choice always brings the work directly beneath it, so there is nothing to look for. Only "What now?" sits above the navigation, and it scrolls away once the work begins (3.7).

Level 1 — the kind of work. Four buttons, always the same four:

Planning • Collection • Analysis • Reporting

Level 2 — where the work happens. Tabs that belong to the current kind of work:

Planning	Collection	Analysis	Reporting
Overview • Settings • Panelists • Theses • AI panelists • section <i>Rounds</i>	(<i>Round 0</i>) • Situation • Invitations • Facilitation • Community • Responses • (<i>Waves</i>)	Summary • Statistics & change	Export • sections <i>Reports</i> and <i>Maintenance</i>

Three newer blocks live inside these: **View as panelist** (the preview) is on the response-link card in Collection, the **Timeline — reasoning and dialogue** is in Collection's Responses, and the pick basket **Thesis material for the next round** is in Analysis above the thesis list — and opens beside the thesis editor in Design when a new thesis is written.

The bottom bar — "Always available". Ten things that belong to no round and no kind of work:

Settings • Responses • Panelists • Theses • Community • Summary • Team & access • Results • Export • Method library

A door carries its destination's name: *Responses* leads to the tab of that name in Collection, *Summary* to the one in Analysis. One place, one name — even when a place has two doors (decision K2).

Two of these deserve a note.

- **Export appears twice on purpose.** Reporting is a *phase* — what you do at the end. Export is a *function* — something you can do at any moment. Same view, two different questions, two different moments.
- **Results is the only item that is not a tab.** It opens a **Results so far** window *over* your current work. Glancing at an earlier round must not take you away from what you are doing, so it is a

window and not a destination. It contains no editing surface at all; there is exactly one way onward — **Open in Analysis** → — and all editing happens there.

Collection's level 2 is a process. The order is the order in which the work is done: *Situation* tells you where things stand, *Invitations* is the first act, *Facilitation* gathers the round's settings (deadline, dialogue phase, visibility and interaction), *Community* is communication — the same view as Community in the bottom bar — and *Responses* is the per-thesis result. *Round 0* comes first when the study runs it, and *Waves* last in the Barometer.

Collection shows the open round. There is no way back into a closed round: its material is read in Analysis, and the server refuses to change a closed round's dialogue, phase rules or schedule. Settings decided before a round begins — opening time and adding a round — are in Planning's *Rounds* section, and exceptions — changing the phase by hand and resetting an old round — in Reporting's *Maintenance*.

The bottom bar also shows your role and AI balance. They are information, not doors: they lead nowhere, but they are always visible.

3.3 Default, not compulsion

The phase decides where you **land** and what gets **weight** — it does not decide where you may go. The rail stays free, because in a Delphi study you return to planning in the middle of collection: round 2's theses are built while round 1 is still running.

When you are looking somewhere other than where the study is, the view says so. It does not stop you and it does not hide anything; it just refuses to let you believe you are looking at the current round when you are not.

3.4 "Applies to:" — the marking that prevents the worst mistake

Planning contains five things, and they are not all the same size:

What	Applies to
Description, phenomenon frame	the study
Settings (orientation, variation, language, AI mode)	the study
Rounds (dialogue, visibility policy, schedule)	each round separately
Theses	each round separately
Panel	the study — but participation is per round

So when you are preparing round 3, you must know whether you are editing round 3's theses or also round 1's. Every planning view carries a marking that says which: **Applies to: the study** / **Applies to: round 3**. Round-scoped editors take the round as an explicit parameter — never "whatever is current".

3.5 AI actions are in the background, behind a control

If the panel includes AI panelists, the Collection view shows a card that says how many there are. The **actions** — generating responses, comments, reactions, revisions — are behind **AI actions ▼**.

The reason is not tidiness. A block of AI operations rendered straight into the main column grows the very column the whole design is trying to shorten, and it puts machine operations in front of the work that is actually the facilitator's at that moment: watching a human panel answer.

If the actions are not available, the card **says why** — the first round has not been opened, round 0's AI sessions are generated from the Round 0 view, or the study has ended. An absent surface that gives no reason is indistinguishable from a broken one.

3.6 One panel, one list

The panel is a single list in which AI participants are marked — `participant_type` is a column on a panelist, not a different kind of object. But a default view listing every panelist would, in several existing studies, be data rather than a panel: measured on 25.8.2026, the database held 80 AI panelists and 33 people.

So the panel view has a filter — **SHOW: All • People • AI** — which is view state and changes nothing. The row always shows all three counts, so what is hidden is readable rather than inferred.

3.7 "What now?" — the card that also says why

The status card's top row names where you are in the arc — *Planning, Round 1 running, Pause after round 1 — prepare round 2, Ending*. Below it are the **next 2–3 steps**, and the order is itself information:

blocker before gap · gap before transition · transition before this phase's work

Each step carries three things. The **reason** says where the prompt comes from (*blocked · missing · next phase · this phase's work*); a blocking reason is amber. The **door** is a button that takes you there and scrolls the control into view. The **rationale** is a small ? that opens the method library at the right article — the rationale is a door, not a sentence written into the card.

The card blocks nothing. When nothing is required, it shows this phase's work: empty is an answer, not an empty card.

A blocker is visible before you hit it. The card asks the gates in advance, so it never prompts you past what would stop you — "Choose the instrument for round 2" appears before you press forward, not afterwards as an error.

The card is at the very top of the screen, above the navigation. You can collapse it when the space is needed for work — but even collapsed, the card shows the first step if that step is a blocker.

3.8 The phase narrows what you see

The Collection view has a row, **Informing and interaction**, below the AI runs: reminder, message to the panel, invitations. It is not a new view but **doors** into existing actions, and its contents follow the phase — in the response phase reminder · message · invitations, in dialogue only message, in the pause message · pause notice.

For the same reason some **sections fold** by phase: what does not belong to the work at hand is a chip rather than a block. Folding is a **default, not a rule** — your own choice wins, and a folded section leaves the chip that opens it.

3.9 The contents row — a view's sections on one line

Views with several sections (Responses, Community, Summary, Statistics & change, Waves) have a **CONTENTS** row at the top. Its chips are the view's sections, and a chip's name is a door: it takes you to the section and opens it if the phase has folded it (3.8). The small × after a chip hides the section. A dashed, faded chip means the section has no content in this study yet — it leads nowhere, because there is nothing there.

For example, Analysis's *Statistics & change* view holds seven analyses — thesis stats, convergence, key figures, dialogicity, shift directions, revisions and refresh — and *Summary* holds, among others, highlights and the epistemic 7×3. The contents row is their index.

3.10 Mail goes out as a letter

Every Pronoia mail goes out in two versions: as plain text and as a letter. The mail program shows the letter; a program that does not display HTML shows the same text as before. You still write only the text — the letter is built around it: the study's name at the top, `[link]` turns into a button where you wrote it, and at the end the contact person and a note that the recipient can reply directly. The letter has no images and no tracking: you will not learn when a panelist opened the message, and it was never promised to them.

4. The same round, from both sides

	The panelist sees	The facilitator sees
Round opens	S1 "your turn": theses, form, description	Collection · Responses: accumulation, reminders, advance
While answering	pulse (process only), own answer confirmed	who has answered, per thesis and per panelist
After own answer	S2, and the distribution for the theses they answered	the whole distribution
Dialogue	S3: others' arguments, anonymous, AI marked	comment stream, AI activity, moderation
Round closes	S4: what came of it, what next	preparing: what the round produced → next round's theses

Read that table downwards and the two designs are one design: **the panelist gets content when it can no longer anchor them; the facilitator gets it immediately, because their job is to see it.**

5. What the interfaces deliberately do not do

Five rules that explain absences you might otherwise read as gaps:

1. **No empty service.** A menu row that opens onto nothing is removed, not left as a promise.
2. **No second computation of the same fact.** There may be several buttons that advance a round; there is one computation of what "advance" means. Two buttons are fine, two answers are not.
3. **AI is marked on every surface where it has an effect** — not disclosed once and then forgotten.
4. **A closed round stays closed.** Its participants and its results are history. Changing them afterwards would silently change a response rate that has already been shown and possibly analysed, and comparison between rounds is the core of the method.
5. **Pronoia describes; it does not interpret.** Distributions, movement between rounds, who said what — those are description. Interpretation is the analysis ecosystem's job, and the boundary is deliberate.

Written 25.8.2026, against the interfaces as rebuilt on 23.–25.8.2026. Button and tab names are taken from the application's own English strings. The specifications behind this text are `docs/UI_Panelistin_nakyma_IA_v2.md` and `docs/UI_Vaihearkkitehtuuri_IA_v2.md`; where they and this guide disagree, they are right and this is stale.

The panelist's process, round by round

This guide walks the panel from the invitation to the final results. It is written for the person who has been invited — *you* is the panelist — and the facilitator who prepares the panel can read it as a description of what their panelists will actually see.

The companion guide, *How Delphi Pronoia's two interfaces are built*, explains the design. This one is the walk: what is on the screen at each point, in order, and why some things are not there yet.

1. What you are joining

A Delphi panel is not a survey. A survey asks you once and adds you up. A Delphi asks you, shows you what the rest of the panel said and why, and asks again — and the interesting part is what happens in between.

Three things follow from that, and they are worth knowing before the first screen:

- **Your reasoning matters more than your number.** The rating locates you on a scale; the reasoning is what other panelists read and what the analysis finally works from. A rating without reasoning is half an answer.
- **You are anonymous to the other panelists.** They see you as a code, not a name: "*You respond anonymously — other panelists see you only as {code}*". The facilitator's visibility into who you are depends on how the study was set up, and the privacy notice says which arrangement applies.
- **Disagreement is not a failure of the panel.** It is often the finding. Nothing in the interface pushes you toward the middle, and the results screens show spread as a fact rather than a problem.

If the panel includes AI panelists, you are told so before you start. "*This panel includes AI panelists*" is stated up front, and every AI contribution is marked wherever it has an effect: beside the argument, and as a hatched share inside the distribution.

2. Before anything opens: the invitation and the consent screen

You arrive by a personal link. It is yours; it identifies you to the study without a password, which is why it should not be forwarded.

The first screen is **Before you begin**. It states what the study is, who is running it, how your data is handled, and — if the panel includes them — that some panelists are AI. You continue with **I have read this and I take part**, and **Read the privacy notice** stays reachable from the menu for the whole study.

This screen is not a formality that the design tries to get past. It is the thing you are agreeing to, so it comes before the content rather than under it.

3. Round 0: the orientation conversation

Most studies begin with **Round 0** — a conversation rather than a form. An AI assistant called **Kastalia** opens the phenomenon with you in your own words before any thesis is rated.

Kastalia says what it is, in its own words, at the start:

"I am an AI, not a person. I help you open up the questions in your own words before the round itself — I don't assess your answers and I won't tell you which view is right. I carry the data; the seeing is yours."

What is on the screen: the phenomenon under study, what the panel design already knows about your angle ("From the panel design, Pronoia knows about you" — correct it in the chat if it is wrong), and a progress line, "Answered {answered}/{total} orientation questions". You write freely; there is no scale.

You end it yourself with **I'm done →**. Then you land on your own home view, and the theses open later.

Round 0 is skipped in some studies. If it is, your first screen is the one in §5.

4. The frame that never changes

From here on, one thing is worth internalising because it makes everything else easy: **the page does not change, its middle does**. The heading is where it was, the menu button is where it was, the form is in the same place every round.

Three parts of the frame are always there:

Part	What it is
Panel pulse	how far the round has got — " <i>{n} have answered</i> ", what is new since your last visit
More services	the index of everything else (§10)
The main column	the only part that changes with the round and the phase

Panel pulse shows process, never content. It tells you that the panel is moving — how many have answered, how many new comments there are — and it deliberately does not tell you *what* the panel thinks until you have answered. There is no target number in it: it says "*27 have answered*", not "27 of 40". A target turns a sign of life into a scoreboard, and a scoreboard tells a late respondent that they are late.

5. Your turn to answer

This is the state you were invited for. The main column holds the theses of the current round, and for each one: the thesis, the scale, and a field for your reasoning — "*Explain your reasoning*".

Answering, in practice

- **The rating and the reasoning travel together.** The placeholder says it plainly: "*Why do you rate this way?*" Some theses are open-ended and have no scale at all; some have two criteria (for example probability and desirability), and then you answer both — they are different questions about the same thesis.
- **Some theses are not scales.** Depending on the study you may be asked to rank options, choose one or several, place a point on a plane, distribute points, mark a time window ("*Earliest*" / "*Latest*", or "*Will never happen*"), or continue a time series. The pattern is the same in all of them: your position, then your reasoning.

- **One at a time, or all at once.** **One at a time** shows a single thesis with "*Thesis {i}/{n}*" and **Previous / Next**; **Show all** returns the full list. This is a reading preference, and it changes nothing else.
- **Your work is kept while you write.** A local draft is saved as you type and the form says "*Draft saved on this device*". If you leave and come back, "*Unsaved draft found*" offers **Restore draft** or **Discard**. The draft lives in your own browser, not in the study.
- **Saving.** Each thesis has **Save**; **Save all ({n}/{total})** submits the ones you have completed. When every thesis has an answer, the view says "*Thank you for your answers!*"

What you do not see yet, and why. During the response phase the panel's *content* — the distribution, the average, other people's arguments — is not shown before you have answered. Not because it is secret, but because seeing it first changes what you write. A first answer that has already been anchored to the panel's average is worth less to the study than a first answer that is yours. The condition is per thesis, not per page: answer this thesis, see this thesis's distribution.

Studies can be stricter or looser about this. The round's visibility setting is one of **Blind**, **After you answer**, or **Open**, and the usual setting for a response phase is *After you answer*.

6. Answered, and the round is still open

Your answer is in. Two things are now true.

You can still change it. "*You have answered all {total} theses. You can return to edit your answers during the round.*" — **Edit answers** reopens the form. Nothing is locked until the round closes — not even when the study moves into the dialogue phase: a thesis you did not get to answer can still be answered during the dialogue.

The theses you answered begin to show the panel. Under a thesis you have answered, the distribution appears: bars for the responses, a line for the mean, a dashed line for the median, a band for the interquartile range, and your own answer marked with ★ and an orange bar. Beside it, the numbers the panel uses — **Md, SD, IQR, n**.

If the panel includes AI panelists, the distribution says so: the AI share is drawn in a second colour **and hatched**, the legend names both groups — **People ({n})** and **AI panelists ({n})** — and n is given split, "*{h} human, {a} AI*". Colour alone is never the carrier of that fact, because colour alone fails for part of the panel.

7. Dialogue: the round where the panel talks to itself

The dialogue phase is what makes this a Delphi. *"You now see other panelists' anonymous answers and arguments. Explore perspectives and consider if you want to refine your own answer."*

Dialogue is open. Here the panel's content does not wait for your answer — during dialogue it is the thing being participated in, and a gate would shut out exactly the panelist who most needs to be let in. (A study can still choose a blind dialogue round; then the interface says so.)

Each thesis card is in one order, and the order is the argument:

1. **the thesis**
2. **your own answer** — so you read the others against what you actually said
3. **the distribution** — where the panel landed
4. **the discussion, behind a link with the numbers in its name** — *"Discussion · {a} arguments · {c} comments"*

The count is in the link on purpose: you can decide whether it is worth opening without opening it.

Inside the discussion

- **OTHERS' REASONING** is the arguments given with the ratings; comments are what the panel said about each other's arguments. They are different things, which is why the link counts them separately.
- Contributions from AI panelists carry the **AI** marker — the same fact as the hatched share in the chart, from the same field.
- If the study uses **role play**, the writer's role is shown next to the handle. A role a human wrote is marked *self-selected*, and a role the facilitator commissioned from an AI *AI role*. You write your own role and place it in a box on your own page, on the **Role play** card.
- On your own reasoning you do not reply but clarify: the button reads **Add a clarification**. Others see the clarification under your reasoning, marked *own clarification*, but it is not counted as dialogue – dialogue is what panelists say about each other's reasoning.
- You can **Reply**, and react with **Insightful**, **Agree** or **Challenge**. A comment that **another panelist** has replied to is locked from editing: *"Locked — the comment has replies"*. Editing it would silently rewrite what someone else answered. Refining your own comment yourself does not lock it. If your position changes after the lock, you say so in a new comment — then the change shows as dialogue, not as history edited after the fact.
- **Translate** renders a contribution in your language when translation is enabled for the panel; **Show original** puts it back.

Changing your mind is the point. If the round allows revision, you edit your answer and say what moved you — the field asks *"Why did you change your mind? (What argument or information influenced you?)"* — and you can classify the change: **More agreement, More disagreement, Nuance / clarification, Reframing.** Then **Save revision.**

Not changing your mind is equally a result. A held position with a reason that has survived the counter-arguments is stronger evidence than a position nobody challenged.

8. Between rounds — the richest screen in the study

When a round closes, the screen does not say "please wait". It says **Round {n} has closed**, and then it shows what the round produced:

"Below is where the panel landed this round: distributions, other people's reasoning, and what was left open. This is the round's interim result — not a waiting screen."

This is the moment the whole method is built around, and it is easy to waste. The round's result is in front of you with no obligation attached: nothing to fill in, nothing due. Whatever view you bring into the next round is formed here.

If the next round has been scheduled you are told when — *"Round {n} opens on {date}"* — and if it has not: *"The next round has not been scheduled yet."* The interim result does not depend on the schedule; time is an addition, not a condition.

In a **Barometer** study this is where a wave ends and the next begins; the interim result is the result, and there is no final screen.

9. When the study ends

"The study has ended" — all rounds are closed. The final distributions, the panel material, the discussion and your own path through it stay readable. If you revised your answers along the way, the interface says how often: *"You made {count} revisions to your answers based on the dialogue."*

What happens after this is outside the panel: interpretation — theming, scenarios, argument and discourse analysis — is done in the analysis ecosystem, and the results reach you as a message rather than as a new screen. Pronoia describes; it does not interpret.

10. More services: the index

More services is an index, not a drawer. It opens a list with one row per service: an icon, a name, one line about what it is, and a mark on the right.

Row	What it answers
About the study	<i>What it's about and where we are now</i>
Messages	<i>Facilitator messages and your replies</i>
My knowledge path	<i>Your contribution and activity</i>
Panel pulse	<i>What the panel thinks right now</i>
Group comparison	<i>How the groups differ</i>
Weak signals & wild cards	<i>Low-threshold foresight — flag an observation</i>
Activity and recognition	<i>Your activity compared with the panel</i>

Four things about this list are deliberate:

1. **A row only exists if the service exists.** Group comparison appears when the study has comparison axes; the signals channel appears when the study collects signals. A row that opens onto nothing is worse than no row.
2. **The mark is a fact, not a nudge.** An unread message is highlighted because it is waiting for you. Everything else is a count, and an empty service says *nothing yet* rather than pretending to be full.
3. **One service at a time, and one step back — Back to services.** Your work stays open behind it: *"Pick a service — your work stays open behind this."*
4. **The privacy and AI notice is always reachable.** It is not a service that has to earn its place by being interesting.

Two of these have a note worth reading.

Panel pulse draws one track per measure: a light band for mean $\pm$ standard deviation, a mark for the mean coloured by *where* it sits on the scale (the colour shows position, not whether the result is good), the group means as dots when the study has groups, and the ends of the scale underneath.

Consensus is not drawn as geometry anywhere: a consensus figure is a classification that needs a calibrated threshold, and that classification belongs to the analysis, not to a bar.

Group comparison is hidden when the groups are too small: *"there are not enough reportable sub-groups of at least 3 people. Results are shown at whole-panel level only."* Small groups identify people. The suppression is the promise working, not a fault.

11. Two ways the same journey is named

A study chooses one of two vocabularies, and the switch reads **View:** with **Story** and **Plain**:

	Story	Plain
Round 0	Threshold · <i>Kastalia</i> · <i>Round 0</i>	Start · <i>Round 0</i> · <i>preparatory chat</i>
The rounds	Lesche · <i>Panel rounds</i>	Rounds · <i>Answering and dialogue</i>
Analysis	Temple · <i>Kassotis</i> · <i>analysis</i>	Analysis · <i>Processing the results</i>

This is naming, not structure. The states, the rules and the surfaces are identical in both; a panel that would find the Delphic imagery strange gets plain words. In the Story view, tapping a station explains what that stage means — *"Tap a station — I'll tell you what it means"*.

Wherever you are, **You are here** marks the current station, and **Round {n} / {total}** says how far along the panel is.

12. If the study is real-time

In a **Real-Time Delphi** there is one round, and answering and dialogue are open at once in it from the start. The orientation conversation (§3) may be part of it too. *"Other panelists' answers appear continuously. You can change your answer at any time."*

The rule from §5 still does the work, thesis by thesis: a thesis shows you the panel as soon as you have answered that thesis. Updates arrive live where the connection allows, and the interface says which mode it is in rather than pretending — *"Updated just now"*, and a tooltip that distinguishes a live connection from fallback polling.

There is no "between rounds" screen in a real-time study, because there is no between. In a **Barometer Delphi**, by contrast, the same theses are returned to one measurement wave at a time: the gap between rounds is the gap between waves, and a thesis may go out of use or be replaced by a new one during it.

13. Your answers, your data

- **Anonymity toward the panel is structural**, not a setting you have to find. Other panelists see a code.
 - **Leaving is possible without erasing the discussion**. You ask the study's contact person, and they carry it out — the promise is kept by someone who is able to keep it. Withdrawal deletes your name, email address and personal link permanently; the answers you have already given stay in the material under the pseudonym, because other panelists have read and argued with them. That asymmetry is deliberate, and it is stated here rather than discovered afterwards.
 - **A closed round stays closed**. Answers to it are not edited afterwards — a response rate that has already been shown, and the comparison between rounds, would otherwise change quietly underneath the study.
 - **The privacy and AI notice** is the authority on how your data is handled in this particular study, and it is one click away from every screen.
-

14. For the facilitator: what is worth saying out loud

Three things panelists ask, that the interface answers but the invitation letter usually does not:

1. **"Do I have to finish in one sitting?"** No. Drafts survive, answers can be edited while the round is open, and the study waits.
 2. **"Why can't I see the results yet?"** Because you have not answered yet — and the moment you answer that thesis, its distribution opens. It is a rule about anchoring, not about access.
 3. **"What happens to what I wrote?"** It is read by the panel under a pseudonym during dialogue, and by the analysis afterwards. The reasoning is the material; the number is the index to it.
-

Written 26.8.2026, against the interfaces as rebuilt on 23.–25.8.2026. Button and tab names are taken from the application's own English strings, and a guard test (`scripts/test_guides_coverage.py`) fails if this guide names a control differently from the program. The

specification behind this text is `docs/UI_Panelistin_nakyma_IA_v2.md` ; where it and this guide disagree, it is right and this is stale.

The facilitator's process, from planning to handover

This guide runs a study from the empty draft to the handed-over result. It follows the order the work actually happens in, and it says at each point what is on the screen, what the panel sees at the same moment, and which actions cannot be taken back.

The companion guide, *How Delphi Pronoia's two interfaces are built*, explains why the facilitator view is organised the way it is. The one-sentence version, because everything below follows from it: **your phase is not the study's phase**. The panelist asks "what is open for me"; you ask "what is my work right now", and while the response phase is open your work is monitoring — and preparing the next round.

0. The map you are working on

Four kinds of work, always the same four buttons:

Planning · Collection · Analysis · Reporting

Ten things that belong to no round and no phase, in the bottom bar marked *Always available*:

Settings · Responses · Panelists · Theses · Community · Summary · Team & access · Results · Export · Share an observation · Method library

The names are the screen's own, not the guide's: a bottom-bar item is a door, and a door is named after the place it leads to.

The phase decides where you **land** and what gets **weight**. It does not decide where you may go: in a Delphi you return to planning in the middle of collection, because round 2's theses are built while round 1 is still running. The button marked *default* is where the study says you belong; when you are somewhere else, the view tells you so — "You are looking at {a}. The study is at {b}." — with a **Back** that returns you. It never hides anything; it only refuses to let you believe you are looking at the current round when you are not.

The navigation is at the top of the screen. The four buttons and the tabs of the chosen kind of work sit in one block that stays in place; a choice brings the work directly beneath it. Collection shows the open round — closed rounds are read in Analysis.

1. Planning: a completeness list, not a wizard

Planning is both a sequence (the first time through) and a place (round 3's theses, built mid-study). Pronoia resolves that by showing the first pass as a list of what is missing rather than as numbered steps — **Building the study:**

"Five things the first round is made of. The order is for reading; readiness is computed from what the study contains."

The five, each of which is also the way in:

Row	What it holds	Applies to
Description	what the study is about, the phenomenon frame	the study
Settings	orientation, variation, language, AI mode	the study
Rounds	dialogue, visibility policy, schedule	each round separately
Panel	who is in it, roles, AI panelists	the study — participation is per round
Theses	what is asked	each round separately

Role play and the anonymous one-shot poll are settings, and each has its own working guide: Role play: a role in a box of the panel matrix and One-shot poll at a seminar: QR code and anonymous answers. Role play needs the panel matrix to be built first.

The list is computed, not ticked off. Readiness comes from what the study contains, and the header states the situation without arithmetic on your part: *"Opening is waiting on {n} item(s)"*, or *"Everything that blocks opening is in place."* Items that are advice rather than obstacles are marked *recommended* — a thesis without a scale is open-ended, which is information, not a fault.

Two of the gates are enforced by the server, not by the screen. A study does not leave draft while the AI mode is unconfirmed, and invitations do not exist until the data-controller declaration is complete.

The gaps are named as they are: "AI mode not confirmed", "controller name", "controller contact", "legal basis for processing", "consent text still has placeholders". The list reads the server's answer instead of guessing it, so you never meet a gate you were not warned about.

And the gap is closed in the view the list leads to. In Settings, under the AI-mode choice, it says which of the two holds: "The AI mode is confirmed" or "The AI mode is not confirmed — a round cannot be opened until it is". **Confirm the AI mode** opens the same choice dialog that otherwise appears only after you hit the gate. The confirmation applies to the selected mode: change the mode and it has to be confirmed again, because a change does not carry the old confirmation with it.

Every planning view says what it applies to. "Applies to: the study" or "Applies to: round {n}", and the panel row says the hybrid out loud: "Applies to: the study · participation is per round". When you are preparing round 3 this is the difference between editing round 3's theses and quietly editing round 1's. Round- scoped editors take the round as an explicit parameter — never "whatever is current".

Rounds are not fixed at creation. "The number of rounds is not locked at creation: the length of a Delphi can be a result rather than a precondition." Dialogue is on by default for an added round; a pure feedback round is made by switching it off, deliberately:

"This round has NO dialogue phase — it collects positions but the panel never argues about them."

Building the panel. The panel is one list in which AI participants are marked — **SHOW: All · People · AI** filters the view and changes nothing. Under it, **PANEL COVERAGE** reports the **Coverage score**, the primary and secondary coverage, the cells that are empty, and — when they are — "Panel supplement recommended before the round".

Reserve enough rounds — an unused round costs nothing. The number of rounds is set when the study is created, and it is not a promise you have to keep: the program never requires a reserved round to be run. You may close the study after round two even if you reserved four; the rounds you did not open simply never existed for the panel. The reverse is harder — a round can be added mid-process, but that is the exception (see the round-management card), not the path. So when the number is in doubt, reserve the larger one.

2. Round 0: orientation, run as a sequence

Round 0 is the one place where the work genuinely is a sequence, and the view shows it as one, and the steps are named on the screen: **Configure, Open round, Send invitations, Monitor chats, Close & go to R1.**

- **Configure** sets the orientation questions, the opening text shown to the panelist, the AI assistant's deepening task and how many follow-ups it may ask. With no configuration yet: *"Create default config"*.
- **Who round 0 is for** is a choice of its own, made before opening. **Whole panel** is the default; **Selected key informants** narrows the round to the ones you tick. The scope covers AI panelists as well as humans, and the same choice drives the AI run, human access to the conversation and the cost estimate. The scope is an explicit choice, not derived from the ticks: one tick must not quietly shut the rest of the panel out.

From then on the scoped number appears everywhere round 0 is discussed — readiness, the R0 counters and the invitation card — beside the panel size. In the collect view the scope has a line of its own, and **Who round 0 is for** → leads back to the choice.

- **Open Round 0** publishes it. From this moment the panelists' links lead somewhere.
- **Invitations** are generated per panelist — **Generate invite links** — and sent from the communication centre or by hand. In a blind study the interface says what it cannot do rather than failing quietly: *"Blind mode hides the addresses: the links are created, but invitations cannot be mailed by hand."*
- **Monitor chats** gives the state of each session — **Started, In progress, Completed, Not opened** — and the average number of turns. You can read any session.
- **AI panelists' orientation chats** are generated here, not from the AI actions control elsewhere: this screen is a lifecycle you run in order, and the two buttons belong to that order. Regeneration deletes the existing AI chats and cannot be undone; the confirmation says so.
- **Close Round 0 → open Round 1** ends orientation and starts the study proper.

A study can skip Round 0 — *"Advance straight to Round 1. Sensible e.g. for a pure AI panel."*

Round 0 belongs to every variation. It comes before the panel opens, and the variation's own character begins only when Round 1 opens — Real-Time Delphi and Barometer Delphi get the orientation round too, and the only way to skip it is the same in all. Before 13.9.2026 the state machine jumped straight to Round 1 in those two, and the card was hidden.

3. Collection: the round-management card

Collection's tabs are in the order of the work: **Situation** • **Invitations** • **Facilitation** • **Community** • **Responses** (with *Round 0* first when it is run). The round-management card is in *Responses*. The settings mentioned below live in their own homes: **visibility & interaction** and the **deadline** in Collection's *Facilitation*, the **opening time** and **adding a round** in Planning's *Rounds* section, and the **phase exception** and **resetting an old round** in Reporting's *Maintenance*.

Everything about running a round happens in **Collection**, and the card marked **Round management** is where the round's state and its controls live: the round, the **Phase**, the accumulation of **Responses**, and a button whose label is the **act** — "*Move to the dialogue phase*", "*Open Round 0 (intake)*", "*Close round 2 (pause)*".

Closing a round is a pause, not an ending. Closing even the last round leaves the study open: rounds can be added. Ending the study has its own control, **End the study** (■), and it is a decision — it asks for a closing letter to the panel and marks the handoff final. Adding a round afterwards undoes it. Before 31 Aug 2026 finality was a side effect of a round, and the facilitator had no way to say *the study is complete*.

There may be several buttons that advance a round; there is exactly one computation of what advancing means, and it happens on the server. That is why the label reads as a destination rather than a verb.

A step says that it was taken. After a successful transition a notice names the destination — "*Round 0 opened · 4 key informants*". Without it, opening looked like it had failed: the only thing that changed was the button's own label, which became the next destination, and a second click in the same spot carried the study past a round.

And a step that leaves something unrun is asked first. If round 0 is unopened or unfinished, moving to round 1 asks for confirmation that names both the destination and what is being left out: "*Opening round 1. Round 0 (intake) will not be run.*" Round 0 is one-shot — orientation gathered afterwards is not orientation but post-hoc justification. Other transitions are not asked: a confirmation asked every time is learned as something to click through.

Before you open a round, two settings are worth a deliberate look, because both change what the panel experiences and neither can be repaired afterwards:

- **Visibility & interaction** — per phase. *Round 1a* — response is **blind by default**; *Round 1b* — *dialogue* is where the panel reads each other. The choices are **Blind**, **After you answer**, **Open**.

- **Deadline** — *not set* until you set it, then "*{n} days left*", **+7 days, Clear**, and *deadline passed* when it has. The deadline is your countdown; the panelist learns of it through a message, not through a clock on their screen.

View as panelist. The response-link card has a button that opens the survey in a new tab exactly as a panelist sees it at this stage — "*The preview identity is not a panel member and nothing is saved.*" The previewer is an identity of its own (an ESIK code), the consent card and the Round 0 conversation are skipped, and every attempt to write returns "*Not saved (preview)*". It does not appear in the response rate or in the reciprocity figures. **End preview** closes the link and removes the previewer; a new one can be opened at any time. This is where the question *what does this look like to a panelist* gets its answer before opening, not after.

4. The response phase: your work is monitoring — and preparing

The study's phase is **response**. Yours is **monitoring**, and the screen is built for watching rather than doing: accumulation per thesis and per panelist, and one intervention that matters — **Remind inactive ({n})**, with a confirmation that names the number, and a result that says what happened: "*Reminder sent to {n} panelists*".

Read the response rate carefully. If the panel is mixed, an AI panelist answers every thesis on every round; a combined rate is therefore partly a property of the panel's construction rather than of its engagement. Pronoia computes the human rate, the AI rate and the combined rate separately and treats the human panel as primary when there are humans — and where there is nobody to divide by, the figure is absent rather than zero, because *nobody was asked* and *nobody answered* are different facts.

The AI panel's actions live behind a control. **AI actions ▼** holds four: *Generate thesis responses*, *Generate dialogue comments*, *Generate reactions*, *Generate revisions*. They are behind a button because a block of machine operations rendered into the main column grows exactly the column this design is shortening, and puts machine work in front of the work that is actually yours: watching a human panel answer.

When the actions are not available the card says why rather than disappearing:

- "*Actions become available once the first round is open.*"
- "*Round 0 is open — the orientation conversations are started here.*"
- "*The study has ended — no new responses are generated.*"

The middle one is not only an explanation. While round 0 is open the card carries the command itself —  **AI panellists' chat**, and **Regenerate AI chats** when there are sessions to redo — next to **Open the Round 0 view**. It used to say the run was started from another view, and there was no way there from that sentence: a sentence pointing at another view is not a path to it.

Every generation passes a cost gate first. **AI run cost estimate** states the scope, the estimate in credits, the budget remaining, who pays and which limit binds — wallet balance, monthly cap, personal cap or this study's own cap — before **Continue and generate**. If the estimate is over: *"The estimate exceeds the remaining budget — some responses may be produced rule-based."*

Meanwhile, the other half of your work. This is when round N+1 gets built. The rail is free, *Applies to:* keeps the two rounds apart, and **Results** lets you glance at an earlier round without leaving what you are doing (§7).

5. The dialogue phase: monitoring and moderation

The panel now reads each other. Your view gains the comment stream, the dialogue threads — *"Dialogue threads (comment-on-comment)"* — and the AI panelists' activity in the same stream, marked.

Dialogue does not close answering. The phase adds a capability — commenting — and removes none: a panelist who did not get to answer can still answer during the dialogue, and changing an answer in the light of someone else's reasoning is the method's idea, not an exception. You need not open anything for it. A comment is locked against editing only once **another panelist** has replied to it; a panelist's own refinement of their own comment does not lock it.

Moderation here is mostly restraint. The dialogue phase is the study's learning phase, and a facilitator who answers arguments becomes a panelist with special powers. What the interface supports is the facilitator's own channel — the **COMMUNICATION CENTRE** — addressed to the **Whole panel**, a **Group**, or an **Individual**, with templates for *Reminder*, *Round open* and *Results* and placeholders ([link] , [handle] , [study]) resolved per recipient. **Group** reaches a panelist through both the home cell and any additional membership on either axis — a person who belongs to two groups receives the group message from both.

Two options there are worth understanding:

- **Also send by email** delivers the message to inboxes as well as the panel.

- **Notice only, no text** sends a mail that says a message is waiting and links to the panelist's own page, leaving the text inside the panel: *"Choose this when the message contains panel content that is not meant for inboxes."*

When a message goes to email it goes out as a **letter**: the study's name at the top, `[link]` as a button and the contact person at the end — the text version travels along for mail programs that do not display the letter. Two more choices:

- **Include »What we ask of you«** adds, after the button, a short three-point description of what the panelist is asked to do. The description is the same in every variation. It is on by default.
- **Preview the letter** shows the letter exactly as it will be sent — the same builder makes both. Placeholders are filled with sample values, and the preview sends nothing.

The delivery report is explicit about who did not get it — recipients without an address, withdrawn panelists skipped, and whether mail was sent at all.

6. The revision phase

Panelists revise their answers in the light of the dialogue, and say what moved them. Your view is the accumulation of revisions and the movement between rounds: **Revisions**, **Revised**, **Revision rate**, under the subtitle *"Panelist changes to R1 responses after dialogue"*.

Revisions is a count, not a share: there is no maximum number of revisions, and the raw figure tells you something about the panel's diligence that a percentage would flatten. **Revision rate**, by contrast, is one share, computed in one place: the *people* who changed their position divided by everyone who responded, humans and AI panelists together. Before 14.9.2026 it was computed in four places in four ways, and the community view showed 1600 %.

7. Glancing at results without leaving your work

Results in the bottom bar is the only item there that is not a tab. It opens **Results so far** over what you are doing: round pills for the rounds that have results, and a row per thesis with n, median and IQR.

The window contains no editing surface at all, and there is exactly one way onward — **Open in Analysis** → — which carries the round you were looking at with it: *"Viewing only — editing happens in Analysis."* When there is nothing to show it says which nothing: *"No results yet — the first round has not been opened."*, or *"Round 0 is the orientation; its results appear in the Round 0 view."*

Glancing is by definition something that does not interrupt. That is why it is a window and not a destination.

8. Between rounds: your richest screen too

The round closes. For the panel this is the interim result; for you it is **preparing**, and it is the only point in the process where the main column sits on the boundary between two kinds of work — analysis produces what planning consumes.

What Analysis gives you here

- **Summary** — distributions, central tendency and spread per thesis.
- **Statistics & change** (thesis stats, convergence, key figures, dialogicity, revisions) and **ROUND-OVER-ROUND MOVEMENT** — *Converged, Diverged, Stable, and "Attention — split theses": "Theses where views diverge the most — often the most interesting part of the analysis."*
- A readability verdict that is about your data rather than your patience: **Not readable yet** (*"Stability cannot be measured before the round is closed and there is something to compare against"*), **Needs action** (*"Coverage is thin. Remind the non-responders before reading the result"*), or **Readable**.

What builds the next round

- **PIRE** — the inter-round engine — proposes, one thesis at a time, how to continue: convergence, stable dissensus, noisy spread, an emergent theme, a bandwagon flag, and an operation to match. You accept, change the operation, flip the mode, or lock the thesis out of the next round. Its own help states the limit plainly: *"PIRE proposes — you confirm. It does not replace deep analysis."*
- **Build the round.** *"Build the next round's survey from the previous round's results — copy as a starting point and edit, or open with the same theses (re-rating)."* The first of these is the default in a round-based Delphi: the next survey is built from the previous results by editing, removing and adding. Plain re-rating is an exception worth justifying — in a barometer it is, conversely, the whole idea. Proposals that involve interpretation are flagged for you to reformulate before opening; you always edit the theses before the round opens. A thesis can also be **copied as the base of a new one** and **hidden** or shown mid-round behind the same gates; hiding does not delete answers.
- **Reading the discussion and picking.** The pause is the moment to go through the previous round's discussion question by question and branch by branch. **Open the whole discussion** shows a thread in full; the **Timeline** shows reasoning and dialogue in one stream, newest first.

Any reasoning or comment can be starred (☆) — *"Mark as thesis material for the next round"* — noting at the same time the **reason for the pick**. Picks collect in the basket **Thesis material for the next round**: *"A pick does not change the comment and is not visible to the panel. It is your memory between rounds."*

- **The basket is source material, not a thesis.** When you write a new thesis, the basket opens beside the editor under **Source material for the new thesis**: *"Write the thesis yourself. The contributions you select are recorded as its sources."* No button turns a pick into a thesis. The selected ones are recorded on the thesis card as **Source material**, which later tells where the question came from. Basket and source material are different things: the basket is your tool between rounds, the source material is bookkeeping. That is why they behave differently when a contribution is deleted — the basket skips it, the source material keeps the record and the reason but the content is emptied.
- **The panel's readiness for the next round** is a different question from the panel's composition: who answered last round, who withdrew, where the coverage gaps are. Composition is edited in Planning; readiness is read here.
- If a round was opened past the pause and you need its survey built after all, **Return Round {n} to preparation** closes it back to the pause and keeps the answers already given.

When you want real interpretation, it happens outside Pronoia: **Send round to DAE analysis** opens the export. Theming, scenarios, argument and discourse analysis and synthesis are the analysis ecosystem's work, under its own quality gates. Pronoia describes; it does not interpret, and this button is where that boundary is crossed on purpose.

9. Reporting and handover

Generate results report produces a descriptive report — summary, per-thesis result cards, the consensus–dissensus map, round comparison R1→R2, and highlights (greatest disagreement, most discussed, most revised). **Print / Save as PDF** and **Download Word (.docx)** take it out of the browser.

The report says what it is, in its own footer, and the sentence is worth quoting to whoever receives it:

"This report is not an analysis but its source material."

Export produces three files for the analysis pipeline: the CSV (quantitative data), the Handoff Package (responses, arguments, dialogue, revision history) and DAE_STATE.json (pipeline control data). Export

sits in two places on purpose: Reporting is a phase — what you do at the end; Export is a function — something you can do at any moment.

Anonymize study is the last step and it is irreversible: panelist names and email addresses are cleared and invite links revoked, leaving anonymous research data — pseudonyms without a key. Do it when the study is over and the material has been exported.

10. Exceptions, and what they cost

Three things sit outside the normal flow. Each is available, each is marked as an exception, and each is recorded.

Situation	What you use	What it costs
The study is stuck in the wrong phase	Change phase (exception)	<i>"The phase normally changes with the advance button. Changing it by hand is an exception: it is logged."</i> The panel sees the change immediately.
You need to look at a closed round	the round selector	<i>"You are viewing closed Round {n} — this view is read-only. Actions always target the current round."</i>
A panelist wants out	withdraw , not delete	Deleting would remove reasoning other panelists have already read and commented on. Withdrawal deletes name, email and link permanently; the answers stay in the material under the pseudonym. Cannot be undone.

11. What the panel sees while you work

You are...	The panel sees
Building the study	nothing — <i>"The study has not started yet. Please come back later."</i>
Running Round 0	the orientation conversation with Kastalia
Opening a round	their turn to answer: theses, form, study description
Monitoring the response phase	process only — " <i>{n} have answered</i> " — plus, per thesis they have answered, that thesis's distribution
Moderating dialogue	other panelists' anonymous reasoning, AI marked, and the discussion behind a link with its counts
Between rounds	the round's interim result in full, and when the next one opens if it is scheduled
Reporting	<i>"The study has ended"</i> — results, material and the discussion stay readable

Read that column downwards and the design is one design: the panelist gets content when it can no longer anchor them; you get it immediately, because your job is to see it.

12. Five things that cannot be undone

1. **Regenerating Round 0 AI chats** — the existing ones are deleted.
2. **Resetting a thesis, a round or a study's responses** — responses, revisions and comments for that scope are deleted. A round reset asks you to type the study name.
3. **Withdrawing a panelist** — name, email and link are permanently deleted.
4. **Anonymizing the study** — identities are removed and links revoked.
5. **Closing a round** — a closed round stays closed. Changing it afterwards would silently change a response rate that has already been shown and possibly analysed, and the comparison between rounds is the method itself.

Item 5 is not only guidance: the server refuses to change a closed round's dialogue, phase rules and schedule. That is why Collection shows only the open round, and a closed round's material is read in Analysis.

Everything else in this guide is reversible, and most of it is reversible by pressing the same button again.

13. What is folded, and when

Some sections are folded by default when they do not belong to the phase you are in. This is a promise, not a surprise: the table below is the whole rule, and it is generated from the same module the program reads (`lib/sectionFolding.mjs`). A guard test fails if the two disagree.

Three things are worth knowing before you read it.

1. **Your own choice always wins.** Open a section by hand and it stays open, whatever the phase says.
2. **A folded section is not a hidden section.** It leaves a chip and a line telling you how many sections the phase folded, so you can open them. In the **CONTENTS** row at the top of the view, a chip's name takes you to the section and opens it.
3. **A section not listed here is never folded.** The rule only removes; it never adds.

`task` is your work in the current phase (build · monitor · prepare · report), `phase` is the round's phase (response · dialogue · revision · closed), and `mode` is the Delphi variation. An empty column means that dimension does not restrict the section.

Section	id	Shown when task is	...and phase is	...and mode is
AI run	run_ai	monitor	—	—
R0 summary	run_r0	monitor	—	—
Respond as AI panelist	run_test	monitor	response revision	—
All responses	run_all	monitor prepare report	—	—
Round movement	ins_movement	prepare report	—	argumentative classical barometer
Epistemic 7×3	ins_epistemic	prepare report	—	—
Refresh	st_load	monitor	—	—
Convergence	st_conv	prepare report	—	—
Dialogicity	st_dialogue	—	dialogue revision closed	—
Shift directions	st_shifts	—	revision closed	—
Revisions	st_revlist	—	revision closed	—
Panel news	com_news	monitor	—	—
Participation	com_participation	monitor prepare	—	—
Signals	com_signals	prepare report	—	—

If a section you needed was folded, that is a fault in this table rather than in your work — and it is worth reporting, because the table is the rule.

14. The line that says where on the arc you are

Above the state card there is one short line. It is the only place that answers the question *where in the whole process am I* — not which screen you have open, but which stage of the study you are standing in.

It says one of these, and nothing else:

The line reads	You are
Design — build the panel and the instrument	before the first round; nothing is running yet
Round 0 in progress — key informant conversations	the orientation round, if the study uses one
Round n in progress — monitor and facilitate	a round is open and answers are arriving
Pause after round n — prepare round $n+1$ from what has accumulated	the round is closed and the next one is not open
Closing — results and next steps	no transition remains; what is left is the report and the letter

Real-Time Delphi and Barometer Delphi speak differently, because they do not walk the same arc (Classic Delphi and Argument Delphi walk it together):

The line reads	Variation
Open round in progress — response and dialogue side by side, no round comparison	Real-Time Delphi
Measurement wave n in progress — the same theses, monitor and facilitate	Barometer Delphi
Between waves (latest wave n) — thesis retirement and preparing the next wave	Barometer Delphi

Why it matters that this is a sentence and not a code. Internally each stage has a short code (A1, B1, A2 ...), and it is in the page's markup where a test can read it. It is deliberately *not* on the screen: a code would be a third vocabulary beside the sentence and the guide, and the codes only exist in round-based Delphi — in the Barometer and in Real-Time Delphi the code is empty, which would make the screen look as though the code were **missing** rather than as though it does not exist.

So the line is written to be read, not decoded. If it disagrees with what you believe the study is doing, believe the line: it is computed from the study's own state, not from the round number.

15. Retirement is the Barometer's way, not a general one

A thesis whose time has passed can be **retired** — but only in the Barometer, and only at the turn of a wave. Everywhere else the program refuses, and the refusal says why: *"Retirement is possible only in the Barometer, between waves. Let the thesis run the round to its end."*

This is a methodological rule, not a limitation of the software.

- **In a round-based Delphi a thesis runs the round to its end.** The whole point of the method is the comparison between rounds: the same thesis, asked again after the panel has seen each other's reasoning. Removing a thesis mid-round would delete the very measurement the round exists to produce.
- **In the Barometer the waves differ only by their date.** The same theses are rated from one time to the next, so a thesis can genuinely reach the end of its usefulness — the future it asked about has arrived, or the question no longer divides anyone. The facilitator decides, and a new thesis can take its place.

What retirement does, exactly:

1. The thesis **stays in the current wave to the end**. Panelists can still rate it, and the wave's results are archived into the wave series as usual.
2. **At the wave turn** it moves to the archive and is no longer asked from the next wave onwards.
3. **Its responses are kept** for analysis. The export marks it `lifecycle = retired`, so the analysis can see that the thesis is no longer being asked — a retired thesis is not a deleted thesis, and the difference matters when someone later wonders why a series stops.
4. **The decision can be cancelled** before the wave turn.

The reason is asked for and stored ("time is up", "the future is visible"). It is worth writing properly: at the turn of the next wave it is the only thing that explains why the series ends there.

Written 26.8.2026, against the interfaces as rebuilt on 23.–25.8.2026. Button and tab names are taken from the application's own English strings, and a guard test (`scripts/test_guides_coverage.py`) fails if this guide names a control differently from the program. The specification behind this text is `docs/UI_Vaihearkkitehtuuri_IA_v2.md` ; where it and this guide disagree, it is right and this is stale.

Role play: a role in a box of the panel matrix

Facilitation (Pronoia)

What role play is for

In an ordinary Delphi the panelist answers as themselves: their expertise is what they bring. In role play the panelist **answers from a role** they write themselves. The role can be their actual position, a future position (*Municipal finance director in 2035*) or a perspective they deliberately choose to represent (*Young rural entrepreneur*).

Role play suits a study when

- you want to **make assumptions visible**: once the role is written out, the reader sees from which position an argument is made
- you are looking at **future actors**: the panelist reasons as a decision-maker of 2035 rather than as today's official
- the panel is **missing perspectives** you want included: a human panelist can take a role, and the facilitator can commission a role from an AI panelist
- the study is used for **learning or workshops**, and changing perspective is the aim itself.

Role play does not suit a study whose value rests on **verified expertise**. A self-selected role is a *claimed* perspective, not a qualification, and it is always marked as such.

Division of work: the facilitator gives the structure, the panelist chooses the box

Role play is built on the **panel matrix**. The facilitator defines the matrix axes and categories, for example the main axis *Expertise* with categories *Technical*, *Economy*, *Administration*, *Experience* and a second axis *Stakeholder*. The categories are **boxes**. The panelist writes their role freely and places it in the box it fits best.

The role is a name, the box is a group. Group comparison, pulse group scores and reports group panelists by box, not by role text. This is the key design choice of role play: if every freely written role were its own group, every group would have one member and $k \geq 3$ protection would hide them all. The boxes are what make role-play results open to structured analysis.

The facilitator's steps

1. **Build the panel matrix** (Panelists → Panel matrix). Choose categories with role play in mind: broad enough that each gets at least three panelists, and clear enough that panelists can choose.
2. **Switch role play on** in the settings and save. The setting explains what the boxes are. If the matrix has no categories, the setting warns you: panelists can then write a role but cannot choose a box.
3. **Tell the panel what a role means in this study** (template below).
4. **Commission AI roles** if you want to supplement the panel. On the **AI panelists** tab, in the add form under **Role play: AI role**, write the role in the **AI role** field and choose a box from the same matrix. The AI answers from its role's perspective.
5. **Watch the boxes fill** before opening the round: a box that stays empty or gets only one or two panelists is not shown as a group.

Template for the panel

In this study you answer from a role. On your own page, on the **Role play** card, write the role you answer from, for example your own position in 2035 or a perspective you want to represent. Then choose the box your role fits best. Other panelists see your role next to your reasoning, but not your name.

What the panelist sees

The panelist's own page has a **Role play** card. The panelist writes their role in **Your role** (at most 120 characters), picks a box on the main axis under **Choose a box** (and optionally on the second axis), and presses **Save role**. The card says: *"Other panelists see your role next to your reasoning and comments."*

In the discussion the role appears next to the handle, marked:

- *self-selected*: the panelist wrote the role
- *AI role*: the facilitator commissioned the role from an AI panelist.

The markers never mix: a role a human chose is never shown as an AI role, or the other way round. When role play is switched off, roles are shown to no one, but the saved role and box are kept.

AI roles as panel supplements

An AI role is useful when the panel lacks a perspective you cannot get from human panelists: young people, future residents, actors from another field. The AI argues from the role's perspective, and its contributions are always marked.

Keep two things in mind:

- **An AI role is a proposal, not a voice.** It brings arguments into the discussion that people can respond to, but it does not tell you what people in that group actually think.
- **The AI counts as a member of the box.** The $k \geq 3$ threshold counts AI panelists in the group size (facilitator decision, 22 September 2026). A box with two humans and one AI is shown as a group. If you want the anonymity of the human group to be certain, keep at least three humans in each box.

Reading the results

- **Compare boxes, not roles.** Group comparison shows how the boxes differ. A single role can be read next to its reasoning, but it is not counted.
- **A self-selected role is a frame, not a sample.** A box's result tells you how the panelists who took that perspective argued, not what a whole profession thinks.
- **Separate the AI share.** In the distribution the AI share is hatched and the counts are split, so you can see how much of a box's result the AI produced.
- **Use roles in interpretation.** The combination of reasoning and role is often the richest material in the study: the same claim made from different roles shows where a view depends on position.

Technical description

- The panelist's role is stored in `self_selected_role` and the boxes in `matrix_primary` and `matrix_secondary`. Call: `POST /api/respond/{invite}/role` with the fields `role`, `matrix_primary`, `matrix_secondary`.
- An unknown box is rejected (`role_play_box_invalid`). When the matrix has categories, a role without a main-axis box is rejected (`role_play_box_required`).
- Adding an AI panelist (`POST /api/studies/{study}/ai-panelists`) takes the same fields; the role is stored as the profile's perspective.

- In the discussion response the role travels in `panelist_role` and `commenter_role` only while role play is on.

See also: Anonymity: identity modes and group reporting.

One-shot poll at a seminar: QR code and anonymous answers

What the one-shot poll is for

The one-shot poll is Pronoia's lightest way to collect answers: no invitations, no sign-in, no handle. The audience scans a QR code with their phones, answers the theses and submits anonymously. Everyone answers once.

The one-shot poll suits

- **taking the temperature of a seminar or workshop:** where the room stands on the theses before and after a discussion
- **the start or end of a Delphi:** a wider audience's view alongside the panel's results
- **teaching:** presenting the method so that everyone gets to try it.

It does not suit a vote whose outcome is binding. Anonymity means no one can verify who answered, and the duplicate-vote guard only works per browser (see **Limits**).

Before the session

1. **Write the theses.** Scale theses are answered with a slider, other theses in text. Short theses about one thing work best in a hall; three to six are enough.
2. **Choose Anonymous one-shot poll in the settings.** Tick **Poll open — accepting responses right now** and save. An unsaved poll will not open.
3. **Check the share card.** A card appears under the tick with the QR code, the **Public link** and the poll's state. The state should say that the code works now.
4. **Test with your own phone.** Scan the code, answer and submit. A second attempt from the same browser says you have already answered; that is correct. Your test answer stays in the results, so allow for it when you read them, or test on a separate copy of the study.
5. **Prepare the presentation view.** **Open presentation view** opens a white full-screen page: the study's name, a large code, the address and the prompt *"Scan the code with your phone and*

answer”. Open it in a browser tab you can project, or take a screenshot for a slide. **Copy link** puts the address on the clipboard if you want to send it to remote participants.

The address carries the study's language. Voters can switch language (Finnish, English, Swedish) at the top of the page.

During the session

What to tell the room:

Scan the code with your phone. Your answer is fully anonymous: no name, no sign-in, and no identifier is stored about you. Answer the theses and press **Submit anonymously**. Everyone answers once.

Keep the presentation view on screen throughout the poll: latecomers can still join, and the address is readable for anyone whose phone will not read the code.

Disruptions:

- **The service is busy.** The voter sees *“The service is busy right now. Wait a moment and submit again – your answers are still in the form.”* Ask people to wait a moment; nothing is lost. The server's limit is sized for a hall (see **Technical description**).
- **The network drops.** The message says the answer was not saved. The answers stay in the form, and submitting works once the connection is back.
- **Someone presses submit without answering.** The page asks them to answer at least one thesis.

Closing the poll

Remove the tick from **Poll open — accepting responses right now** and save. After that the code opens a page that reads *“The poll is closed.”* Anyone who submits at the very moment of closing sees a message that the poll was closed before their answer arrived.

You can reopen the poll the same way, for example for a second measurement after a discussion. Note that the same-browser guard applies to the whole poll: to measure the same audience twice, run the second measurement as a separate study.

Results

The answers go into the round's results and distribution like any other answers. Each submission is stored as an anonymous participant who does not appear on the panel roster or in group comparison. That is why one-shot poll results cannot be broken down by group: voters have no box in the panel matrix.

Read the result like this:

- **The distribution shows where the room stands**, not the expert panel's position. If you set it beside Delphi results, say in the report that it comes from a different group.
- **The reasons are anonymous and one-off**. They cannot be clarified or answered, so they are material, not dialogue.
- **The number of respondents is the number of submissions**. Without an identifier, it does not tell you how many different people answered.

Limits

- **The duplicate-vote guard is per browser**. The same browser answers once. The guard stops accidental repeats, but not deliberate ones: another browser or a private window can answer again.
- **No commenting**. Voters can write a reason, but cannot comment on others' answers.
- **No follow-up**. A one-shot poll creates no own page, no revision and no next round. To engage participants further, invite them to the panel in the usual way.

Technical description

- The poll page is `/poll/{study}?lang={language}` and the presentation view `/poll/{study}/qr`. The QR code is drawn in the browser (error correction level M, four-module quiet zone); the address never passes through an outside service.
- `GET /api/poll/{study}` returns the theses of an open poll and `POST /api/poll/{study}` stores the answers. Each submission creates an anonymous participant (`participant_type = anonymous`).
- The server rejects a closed poll (403) and overload (429); the page tells the voter about both.
- The per-address limit is 1,500 submissions and 3,000 page loads in ten minutes. Phones in a hall often reach the server from the same address, so the limit is sized for a hall, not for a single user.

- The duplicate-vote guard lives in the browser's storage; the server stores no identifier for the voter.

See also: Anonymity: identity modes and group reporting.

Taking part in the open Delphi study

Participation (Pronoia)

Metodix always has one open Delphi study running, and anyone interested in the topic can ask for an invitation. This guide tells you in two minutes what you are joining. Answering itself takes about fifteen minutes.

No account needed

A personal link arrived in your email, and that is all you need. There is no password and nothing to register. The same link brings you back to the panel round after round, so keep the message.

Please don't pass the link on. It is your seat in the panel, and if someone else answers with it, their views are recorded as yours.

A Delphi study is not a survey

A survey adds opinions up. A Delphi study reads them. A round goes like this:

1. **You answer first.** You rate the theses before you see what anyone else said. The order is deliberate: the first number you saw would stick to your own.
2. **You give your reasoning.** The reasoning is the most important part of your answer. The rating says where you landed; the reasoning says why, and that is what the analysis reads.
3. **You read the others' reasoning.** Under **OTHERS' REASONING** you see the rest of the panel anonymously. You can comment on a single piece of reasoning and ask its author for more.
4. **You may revise.** If some argument makes you think differently, change your answer and say what changed it. Changing your mind is not weakness but a result of the method, and it appears in the analysis as a finding of its own.

The aim is not unanimity. The aim is that those who disagree know why they disagree.

Anonymity is structural, not a promise

The other panelists see only a pseudonym. They see your rating and your reasoning; your name, your organisation and your email address stay out of sight. This is not politeness but a condition of the method: anonymous reasoning is judged by its content rather than by its author's standing.

In the distribution chart your own answer is marked **you**, and in the discussion your own turns appear to you as **You**. To everyone else the same row sits under a pseudonym.

Only the organisers at Metodix see your contact details, and they use them solely to send the invitation and the round messages.

What else is on the page

- **DISCUSSION** gathers the conversation around a thesis into a thread, and **YOUR REASONING** shows your own text, which you can refine.
- **More services** shows what has been stored about you, and it is also where you can withdraw from the study.
- **Weak signals & wild cards** is a place for observations that do not fit the theses but may still change the picture of what is coming.

If something doesn't work

Open **Share an observation** and write down what you tried and what happened. The message goes to the people who build Pronoia, not to the study's facilitator, and it has no effect on your answers. The most useful thing you can tell us is where you looked first — even if it was the wrong place.

Short answers

Question	Answer
Do I need an account?	No. The personal link is enough.
How long does this take?	About fifteen minutes per round. You can answer at your own pace.
Will anyone see my name?	Not the other panelists. The organiser sees your contact details for the invitation.
Can I change my answer?	Yes, for as long as the round is open. The change and its reason are part of the material.
Can I stop halfway?	Yes. Withdrawal is in the More services menu.

The facilitator trial in five stages

Facilitation (Pronoia)

You have a trial of Delphi Pronoia. In ten minutes this guide walks you through what running a Delphi study actually involves.

The stages are written as goals: each one says what you are trying to find out, not which button to press. Finding the route is part of the trial — and if you can't find it, that is the most valuable thing you can tell us.

1. Open your own example study

The invitation link lets you create an account and a password. After that your desk holds a study that already has a panel, theses and answers. It is yours: you can change anything in it without breaking anything.

Goal: find out what the study asks and how many people have answered.

Both answers are on the **Theses** and **Panelists** tabs. Note the basics of the method while you are there: a thesis is a claim to take a position on, not a question.

2. See where the panel landed

Goal: find one thesis the panel disagrees on, and find out on what grounds.

The distribution shows how the answers fell. Below it are the statistics: mean, median, standard deviation and the middle half. In a Delphi study the spread is often more interesting than the mean, because it tells you whether the panel is of one mind or split.

The reasoning, though, is why Delphi studies are run at all. Read it and see whether the disagreement lies in the ratings or only in the reasons behind them.

3. Read the dialogue

Goal: find a place where someone revised their position, and work out what changed it.

Panelists comment on each other's reasoning and may revise their own rating. The revision and its reason are stored separately. This is what separates a Delphi study from a survey: the material keeps a record of which argument moved whom.

4. Change something

Goal: add one thesis of your own and see how it looks to a panelist.

Use the panelist preview: you see the same view as your panel, and nobody is notified. It is the fastest way to check whether a thesis makes sense.

Keep one rule in mind: a thesis added while a round is open means that part of the panel answered a different set from the rest. Pronoia does not prevent it, but it does record it.

5. Build a study of your own

Goal: take one real question of yours as far as the theses.

The wizard on your desk asks for the variation (Argument, Classic, Real-Time or Barometer), the number of rounds and the default scale. On the **AI panelists** tab you can add synthetic perspectives if you want to see the dialogue move without a real panel.

When you invite real panelists, the **Invitations** tab creates a personal link for each of them. You can send them from Pronoia or copy them into your own email.

The limits of the trial

Limit	Amount	Why
Duration	14 days	A trial is for getting to know the tool, not for production use.
AI	5 euros' worth	AI panelists and analysis helpers are charged by use.
Real panelists	10	Across all of your trial's studies together. AI panelists you may add freely.

Hitting a limit is not an error but the end of the trial. Tell us what you were doing, and we will agree on what comes next.

Tell us what didn't work

At the bottom of the screen there is **Share an observation**. Use it freely. The valuable part is not that something was broken, but where you looked first and what you expected to happen. Pronoia's most common defect is a feature that exists but cannot be reached from the screen where it is needed — and only a new user can see that.