

USER GUIDE

Panel design

Delphi ecosystem — user guides

English

Version 1.0 · 2026-09-23

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

Contents

PANEL DESIGN

Building the Delphi panel	3
The panel as a designed epistemic space	3
The two-axis matrix is primary	3
Orientation sets the panel type	4
Perspective: position + view (Hautamäki)	5
AI panels (Route B / C)	5
Reading the arguments: the epistemic lens	6
Panel diagnostics (coverage, PDI, lock-in, PHS)	6
Quality gates before collection begins	7
Building the panel in practice	7

Building the Delphi panel

The panel defines the **entire argument space** of a study: every argument is one panel member's response to one thesis, so the arguments a study can ever produce are *theses × panel members*. Design the panel well and its coverage — and its blind spots — are known **before any data exists**. This guide is the methodological centre of planning a Delphi study in the ecosystem.

Planning (DPE) chooses orientation and drafts the panel. **Facilitation (Pronoia)** is where the panel is finalised, validated and gated before collection. The panel is a **designed epistemic space**, not a recruitment convenience.

The panel as a designed epistemic space

A panel is not a random set of experts; its reach and gaps are deliberate and documented. The guiding principles the facilitator works with:

A **homogeneous panel** (same education, sector, culture) is an epistemic risk and is recorded as such — even if every member is individually excellent. The target is **commensurability** — enough shared ground for genuine dialogue — not full consensus. **Disagreement that is well covered is a feature, not a defect**. And watch for **lock-in**: a single perspective claiming universality, which can masquerade as dialogue simply because it is loud.

The two-axis matrix is primary

The panel's primary structure is a **two-axis matrix**. Its cells hold the **roles** that panelist behavior is grounded in. The **facilitator ultimately defines the axes and their classification** — the matrix is the design surface, and the templates below are only starting points.

The axes are **context-dependent**. The canonical choices:

Panel type	Axis 1 (primary)	Axis 2 (secondary)
Expert panel (N/ S/R/D)	Expertise (research / practice / administration)	Stakeholder affiliation (public / private / civil society)
Participatory panel (E)	Concerned-party position (resident / civil servant / decision-maker / NGO)	Expertise / theme area
Corporate strategy	Concerned-party group (leadership / staff / customers / owners)	Competence area (strategy / finance / operations / market)

Two properties matter in practice. A **cell may hold 0..n panelists** — several voices in one role, or none. **Empty cells are expected**, and a systematically empty cell is a **structural blind spot** (not a random gap). Naming blind spots before round 1 is part of designing the space; they become candidates for supplementation later.

Orientation sets the panel type

The study orientation chosen in planning (N, S, R, E, D or D+) decides whether the panel is built as an **expert matrix** or an **observer matrix**, and which axes, selection logic and minimums apply.

Orientation	Panel type	Selection basis	Practical requirement
N / S / R / D (intentional)	Expert matrix	Expertise, stakeholder representation, competence domain	At least 3×3; ≥3 stakeholder groups; ≥2 domains; power balance considered
E (emergent)	Observer matrix	Observation distance (inside / edge / outside), not expertise	At least 2×3; ≥3 observation distances; ≥3 "outside" roles; no decision-makers; each role names its position and "what it sees"
D+ (four-path)	Two parallel panels	Intentional and emergent in parallel	Both requirement sets apply independently

Emergent studies are about *where you stand to watch*, not credentials — which is why decision-makers are excluded and observation distance replaces expertise as the selection basis.

Perspective: position + view (Hautamäki)

The optional but recommended **perspective layer** rests on a simple idea: **perspective = position + view**. The same phenomenon genuinely looks different from different positions — not because anyone is wrong, but because the phenomenon is multidimensional. The aim is to identify what each position *can* see and what it *necessarily cannot*.

Following Hautamäki's viewpoint relativism, the deeper point is the **multiplicity of perspectives in general** — a panel should span the perspective space, not just tick boxes. The familiar four archetypes (constructivist optimist, empirical realist, critical-conservative, eco-social holist) are **one useful special case**, not the whole space; add or vary perspectives as the phenomenon demands.

When the layer is active, each role carries three attributes that make its standpoint explicit:

Attribute	What it captures	Values
Observation position	Where the member stands relative to the phenomenon	inside / edge / outside / below / above
Epistemic type	The kind of knowledge the member brings	formal / experiential / tacit
Value basis	The standpoint's underlying values	short free text

The perspective layer **complements** the matrix cell: the cell sets *what* a member speaks to (its role); the perspective sets *how* it reasons.

AI panels (Route B / C)

A study can run wholly or partly with **AI panelists** — a full-AI perspective panel makes an end-to-end run possible without recruiting humans, which is valuable for testing the whole process and for **human-validated AI-analysis development** (AI proposes → human validates → κ/calibration → improve).

An AI panelist sits in a **matrix cell** (its role) and carries an explicit **perspective profile** that complements it. Because the AI acts on that description, it must be **precise and sufficient** — this is enforced by the **AI Panelist Description Standard** (docs/ ai_panelist_description_standard.md): each AI panelist needs a filled cell (expertise × stakeholder), a distinct perspective label, the four Hautamäki dimensions (ontological / epistemological / axiological / temporal) as real descriptions, `perspective_axes` in [0,1], a bias

direction and the model used; across the panel the labels are unique and the perspectives diverse. Route B/C also requires a GenAI disclosure, and humans stay in the loop: every AI response is reviewed before a round closes.

Reading the arguments: the epistemic lens

As arguments arrive, Pronoia reads each one on two dimensions — **where the claim draws from** (its knowledge base) and **what it does to the thesis** (its function).

Knowledge base — seven categories: **EMP** empirical (data, measurement) · **KOK** experiential (personal/professional, tacit) · **NOR** normative (values, ethics) · **TEO** theoretical (models, frameworks) · **INT** interest-based (stakeholder/strategic position) · **EPV** uncertainty (acknowledged ignorance) · **IMA** imaginative (vision, radical reframing).

Function — three categories: **TUK** supporting (reinforces the thesis) · **HAS** challenging (questions, weakens, refutes) · **AVA** opening (introduces a new viewpoint, widens the frame).

Seven knowledge bases × three functions give a **21-cell map** of the argument space. Empty cells are the panel's blind spots — viewpoints that never appeared. The **Epistemic Coverage Index** is how many of the 21 cells are filled; a typical run fills roughly 9–14 of 21, and very low coverage signals a need to add roles that bring the missing viewpoints.

Panel diagnostics (coverage, PDI, lock-in, PHS)

Description quality is *per panelist*; **coverage** is a property of the whole panel. The ecosystem computes it deterministically — the computable core of **MaxPerspective (M5)**, aligned with Hautamäki's viewpoint relativism, available live in the new-study wizard and as a tool over any panel:

- **Matrix coverage** — filled vs empty cells (Absent cell = structural blind spot) and fill rate.
- **Dominance & diversity per axis** — aspect dominance **AD%** per category, the dominance class (DOMINANT >50% / STRONG / MODERATE / WEAK / TRACE / ABSENT), and the **Perspective Diversity Index $PDI = 1 - \sum(AD_i/100)^2$** (Herfindahl). **Lock-in** is flagged when one category exceeds 50%.
- **Structural PHS (0–50)** — $PDI \times 25 + (1 - \max AD/100) \times 25$. The full Perspective Health Score also weighs *awareness* and *openness*, which are qualitative (the six-phase MaxPerspective analysis) and not inferable from a panel spec.
- **AI perspective spread** — over `perspective_axes`, the spread of the AI panel across the perspective space; low spread is perspective lock-in even when the matrix looks covered.

- **Supplement recommendations** — which categories to fill (Absent blind spots), where to reduce concentration, and how to broaden perspective spread.

As a fast human-facing summary the facilitator may also record a **PHS quick score (0–12)** from four diversity dimensions — position diversity, epistemic diversity, value-basis variation and marginal voices, each 0–3: **0–5** diversity too thin (supplementation required) • **6–8** adequate with gaps (Hautamäki perspective theses recommended) • **9–12** excellent diversity.

Quality gates before collection begins

Before a panel is accepted it must pass orientation-specific gates that turn the principles above into checkable minimums:

Check	Intentional (N/S/R/D)	Emergent (E)
Matrix size	≥ 3×3	≥ 2×3
Stakeholder groups / observation distances	≥ 3 groups	≥ 3 distances
Domains / outside roles	≥ 2 domains	≥ 3 outside roles
Decision-makers	Power balance noted	Exactly 0
Per-role attribute	Expertise named	Observation position + "what it sees" named
Theses	≥ 3	≥ 3

If a gate fails, the panel is not yet ready — the facilitator adjusts composition or theses and re-checks. The gates keep coverage and balance honest before any arguments are gathered.

Building the panel in practice

Ready-to-run **templates** give a validated starting point for each panel type — Classic foresight, Argument Delphi, Barometer, Participatory civic, Real-Time, Corporate strategy and a full-AI perspective panel — each a contract-conformant package you instantiate in one step (from the study wizard or by importing a planning package) and then edit.

A working sequence for the facilitator: pick a template (or start blank) → set the two axes for your phenomenon → place roles in the cells (leaving deliberate blanks) → for AI panels, write each perspective profile to the standard → read the live diagnostics (PDI, lock-in, blind spots, structural PHS, AI spread) and act on the recommendations → pass the quality gates → begin collection.

The panel is where a Delphi study's epistemic reach is decided. Everything downstream — the arguments, the analysis, the scenarios — is bounded by the space you design here.