In-Flight Insights: Modified Delphi for Corporate Foresight 2035

Corporate R&D organisations must anticipate long-horizon challenges that demand basic research, yet corporate foresight outputs often arrive too late to shape strategic conversations. This is particularly problematic when foresight is expected to inform near-term agenda-setting while addressing a distant horizon. Prior research highlights that foresight value depends on how tightly insights couple with ongoing strategic decision processes, and that weak integration often undermines impact (Rohrbeck & Schwarz, 2013; Vecchiato, 2012). This paper reports on a **modified Delphi panel** conducted in 2025 within the R&D unit of a large telecommunications company to inform a **10-year horizon to 2035**. The paper addresses two linked challenges: (1) how to design a Delphi process that combines depth and interdisciplinarity with corporate time-to-use constraints, and (2) how to avoid the common trap of separating data collection from analysis and utilisation, which delays organisational impact.

Theoretical Positioning and Research Gap

Corporate foresight research has identified a persistent tension between methodological rigour and organisational responsiveness. Rohrbeck and Schwarz (2013) demonstrate that foresight impact depends less on the quality of insights produced than on the degree to which those insights are integrated into ongoing strategic processes. Similarly, Vecchiato (2012) argues that foresight initiatives frequently fail not because of weak analysis, but because of weak **coupling mechanisms** between the foresight function and decision-making routines. This decoupling is compounded in Delphi-based studies, where the sequential logic of rounds — collect, synthesise, report — structurally separates knowledge production from knowledge use (Rowe & Wright, 1999).

Adaptations of the Delphi method in organisational contexts have addressed panel composition (Hasson & Keeney, 2011), consensus thresholds (von der Gracht, 2012), and hybrid qualitative-quantitative designs (Okoli & Pawlowski, 2004), but have not systematically addressed the **temporal integration problem**: how to ensure that foresight outputs become actionable during, rather than after, the research process. This paper conceptualises that problem as the central design challenge and advances a methodological response grounded in real-time strategy formation (Eisenhardt & Martin, 2000) and organisational learning theory (Argote, 2013).

Method

The panel comprised **three rounds**. **Round 1** replaced the conventional Delphi survey with **90-minute qualitative expert interviews**, designed as an exploratory phase that surfaced future-facing speculations as well as present-day "expert realities" (assumptions, constraints, emerging tensions) from top specialists. **Rounds 2 and 3** followed a more classical Delphi logic using structured questionnaires, iterative feedback, and refinement of statements. The panel intentionally mixed telecommunications experts with contributors from other fields (including climatology, economics, sociology, anthropology, history, and philosophy) as well as specialists in adjacent domains relevant to telecom R&D (e.g., quantum technologies). In Round 1, **37 experts** participated; **24** continued through Rounds 2 and 3.

Interview transcripts were analysed using thematic coding (Braun & Clarke, 2006), with an initial open-coding pass followed by axial coding to identify clusters of forward-looking claims, present-day constraints, and expert assumptions. Statement development proceeded through structured distillation: coded segments were reformulated into falsifiable propositions by a two-person research team working independently, with inter-rater reliability checked at the statement-construction stage. Bias mitigation measures included cross-disciplinary panel composition, anonymisation of attribution during survey rounds, and explicit probing of outlier positions in Round 3 facilitation. Consen-

sus was defined as **≥75% agreement**, with moderate and high consensus thresholds at **≥85%** respectively, consistent with established practice in corporate Delphi applications (von der Gracht, 2012). From **100+ statements**, **55** reached consensus.

Methodologically, the paper advances a "**harvest now, use it now"** operating model for corporate foresight: analysis and use are treated as part of the process rather than a post-hoc phase. Instead of waiting for a final report, the team produced **interim syntheses** and **learning artefacts** while the panel was still ongoing, without abandoning the iterative logic of Delphi. This model instantiates a form of **concurrent validity** in organisational terms: insights are tested for decision-relevance in situ, and the foresight team adapts framing and emphasis in response to emerging stakeholder questions. Organisationally, the approach strengthens future orientation across R&D and a wider stakeholder community while managing common corporate constraints such as limited attention, competing priorities, and report fatigue.

Conceptual Contribution: The "Harvest Now, Use It Now" Framework

The "harvest now, use it now" model is not merely an operational convenience; it constitutes a theoretically coherent departure from conventional Delphi logic. We propose it as a **transferable framework** with three constitutive components and defined boundary conditions.

- **Component 1 — Staged artefact production.** Rather than a single terminal report, the process generates a layered sequence of artefacts (interview synthesis briefs, mid-panel signal maps, Round 2 consensus dashboards) calibrated to organisational decision cycles. Each artefact is designed for a specific audience and decision moment, drawing on Weick's (1995) notion of sensegiving as an active process through which foresight teams shape how futures are interpreted.
- **Component 2 — Parallel utilisation tracks.** Insight use is scheduled in parallel with insight generation, not sequentially after it. This requires explicit **use-planning** at the outset: identifying which stakeholders need what kind of input at which moment in organisational planning cycles. This mirrors Mintzberg's (1987) argument that strategy forms through doing, not only through planning.
- **Component 3 — Iterative organisational closure.** The process is designed to reach **partial closure** at multiple points rather than a single definitive endpoint, allowing the organisation to act on available evidence without waiting for full consensus. This corresponds to satisficing logic (Simon, 1955) adapted to foresight contexts.

Boundary conditions. The framework is most applicable in corporate R&D contexts characterised by: (a) significant time pressure combined with methodological credibility requirements; (b) distributed stakeholder communities with heterogeneous information needs; and (c) foresight horizons of 7–15 years, where radical uncertainty makes premature closure costly. It is less appropriate in regulatory or policy foresight contexts where sequential public consultation is legally mandated.

Selected Empirical Insights

Three clusters of findings illustrate the substantive character of the panel's consensus and the analytical value added by the cross-disciplinary design.

Trust as a scarce structural resource. The panel's most consequential convergence concerned not a specific technology but a systemic condition. Across disciplines, experts converged on the proposition that by 2035, trust will have ceased to function as a social default and will instead become something that must be actively produced, technically verified, and institutionally defended. Two high-consensus statements anchor this cluster: 81% of experts agreed that the proliferation of AI-generated content and synthetic identities will erode shared notions of authenticity, making technological verification systems essential for maintaining social trust and institutional credibility; and 80% agreed that deepfake technologies will extend beyond audio and video to reliably mimic indi-

vidual behavioural signatures — writing style, decision-making patterns, biometric patterns, and interaction habits — making synthetic identity replication extremely difficult to detect without advanced authentication systems. These findings emerged from the convergence of telecommunications engineers, anthropologists, sociologists, and legal scholars, and reframed what had initially been treated by technical experts as an engineering problem into a governance problem with no clear institutional path to resolution. For telecommunications operators, the panel identified a strategic corollary: the ability to provide a verified trust layer — confirming not only message delivery but source identity and authorisation chain — may become one of the sector's most significant capabilities.

Zoned autonomy and swarm-level systemic risk in multi-agent systems. The panel reached 84% consensus that AI agents performing high-stakes or strategic tasks will still require continuous human oversight by 2035, while simultaneously reaching 83% consensus that agents will routinely operate in multi-agent ecosystems — negotiating, coordinating, and forming short-lived coalitions — at machine speed. This dual structure — tightly fenced where accountability is legally mandated, effectively ungoverned where speed takes priority — produced the panel's most analytically novel contribution: the identification of swarm-level emergence as the primary governance failure mode. Experts across computer science, economics, and complexity theory converged on four failure patterns of particular concern — cascade failures, silent convergence, brittle dependencies, and specification gaming at scale — none of which are visible at the level of individual agent behaviour. This finding was not present in Round 1 expert framings, which focused on individual agent capabilities; it emerged through the iterative consensus process as cross-disciplinary interaction surfaced the sociology and ethology of agent populations rather than their technical properties. The panel reached 76% consensus that machine-to-machine interactions will dominate critical economic and infrastructural processes by 2035, operating at speeds increasingly beyond real-time human intervention.

Energy as a structural constraint on digital infrastructure. With 84% consensus — among the highest in the study — the panel expects leading technology firms to pursue substantial energy self-provisioning by 2035: private solar and wind assets, microgrids, battery storage, and long-term power agreements that reduce dependence on public grids. This is not primarily a forecast; it is a description of a bifurcation already underway. A further 76% agreed that rising computational demand driven by AI, simulation, and robotics will outpace electrical grid expansion in several regions, making energy availability a structural constraint on digital infrastructure growth. The strategic implication — identified through the convergence of climatologists, economists, and telecommunications engineers, rather than from telco-internal analysis — is that energy resilience has ceased to be a secondary operational concern and has become a core architectural requirement and a potential axis of competitive differentiation. For telecommunications operators, the panel identified an asymmetric risk: operators may be required to finance the dense, energy-resilient infrastructure that the digital ecosystem requires, while hyperscalers integrating across energy and compute capture the most profitable orchestration layers above the network.

Across all three clusters, the cross-disciplinary panel design was the generative condition for findings that would not have emerged from a single-discipline or intra-sector study. The "harvest now, use it now" model was applied within each cluster: interim synthesis briefs were produced after Round 1, surfacing these cross-disciplinary reframings for stakeholder use before full consensus had been established in Rounds 2 and 3.

Benefit

The paper offers a transferable playbook for running a modified Delphi in corporate R&D settings where **time-to-insight** is as critical as methodological credibility. Participants will gain:

- design principles for three-round Delphi processes that begin with qualitative interviews and transition into survey-based convergence, including analytical procedures for statement development and inter-rater reliability management;
- a staged workflow that integrates data collection, analysis, and utilisation ("harvest now, use it now") while preserving the integrity of iterative consensus-building, with defined roles for use-planning alongside data collection; and
- stakeholder practices that increase uptake by turning exploration into organisational learning rather than a static deliverable, illustrated through the framework's three components and its boundary conditions.

The paper will illustrate the approach through **three anonymised "in-flight use" patterns**: how interim synthesis informed expert communication and internal sensemaking during Round 2; how the energy-constraint finding was mobilised to reframe investment discussions before Round 3 was complete; and how evolving outputs were made decision-relevant to a 2035 strategic planning cycle without premature closure on unresolved expert disagreements.

Key References

Argote, L. (2013). Organizational Learning: Creating, Retaining and Transferring Knowledge (2nd ed.). Springer.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77–101.

Eisenhardt, K. M., & Martin, J. A. (2000). Dynamic capabilities: What are they? Strategic Management Journal, 21(10–11), 1105–1121.

Hasson, F., & Keeney, S. (2011). Enhancing rigour in the Delphi technique research. Technological Forecasting and Social Change, 78(9), 1695–1704.

Mintzberg, H. (1987). Crafting strategy. Harvard Business Review, 65(4), 66–75.

Okoli, C., & Pawlowski, S. D. (2004). The Delphi method as a research tool: An example, design considerations and applications. Information & Management, 42(1), 15–29.

Rohrbeck, R., & Schwarz, J. O. (2013). The value contribution of strategic foresight: Insights from an empirical study of large European companies. Technological Forecasting and Social Change, 80(8), 1593–1606.

Rowe, G., & Wright, G. (1999). The Delphi technique as a forecasting tool: Issues and analysis. International Journal of Forecasting, 15(4), 353–375.

Simon, H. A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics, 69(1), 99–118.

Vecchiato, R. (2012). Environmental uncertainty, foresight and strategic decision making: An integrated study. Technological Forecasting and Social Change, 79(3), 436–447.

von der Gracht, H. A. (2012). Consensus measurement in Delphi studies: Review and implications for future quality assurance. Technological Forecasting and Social Change, 79(8), 1525–1536.

Weick, K. E. (1995). Sensemaking in Organizations. Sage.