



University of St.Gallen

Center for Geopolitics and Corporate Strategy

WHITE PAPER SERIES

Issue 3 · September 2026

Decision Intelligence for a Geopolitical Age

How Delphi Studies, Open-Source
Intelligence, and LLM-Enabled
Strategy Workflows Help Firms
Respond to Geopolitical Disruption

Roger Moser

From insight to impact.

Moser, R. (2026). *Decision Intelligence for a Geopolitical Age
How Delphi Studies, Open-Source Intelligence, and LLM-Enabled
Strategy Workflows Help Firms Respond to Geopolitical Disruption.*
St. Gallen Center for Geopolitics and Corporate Strategy, White Paper
Series, Issue 3, (September 2026)

Roger Moser, Prof. Dr.

Titular Professor at the University of St.Gallen (HSG) and faculty at Macquarie Business School (AUS) and IIM Udaipur (India).

- His research spans decision intelligence, supply (chain) management, and international business strategy, with a particular focus on India and emerging markets.
- He has worked with organisations including Daimler Trucks, BMW, Robert Bosch, Airbus, SatSure, Bartholet and UK Trade & Investment, among others, on long-horizon strategic foresight and structured expert panels.



Executive Summary

Between January and March 2025, a sequence of trade-policy announcements out of Washington rewrote the operating assumptions of several hundred European manufacturers in under ninety days. Most could not respond in time, not because the relevant expertise did not exist, but because it sat immobilised in the heads of people who were never asked a structured question. This is the defining problem of the geopolitical age: disruption now moves at the speed of the news cycle, while the machinery firms use to make sense of it still moves at the speed of the planning cycle.

This white paper argues that geopolitical adaptation is not, in the first instance, a matter of better information. It is a matter of better institutional capability for judgment and decision-making. The expertise needed to read a geopolitical event already exists; what most firms lack is the infrastructure to mobilise, aggregate, and translate it quickly enough to act.

It describes one such infrastructure, a Decision Intelligence Platform developed at the University of St. Gallen, built from three integrated layers: open-source intelligence (OSINT) for signal detection, online Delphi expert panels for structured judgment, and LLM (Large Language Model)-enabled workflows that translate that judgment into concrete consequences for a specific firm. Machines scan, detect, and decompose; humans judge, contextualise, and disagree productively. The result is an assessment fast enough to matter and auditable enough to defend in a boardroom or a regulatory filing.

For executives, five implications follow:

- Build intelligence infrastructure before you need it: The speed of crisis response depends entirely on expert registries, frameworks, and baselines assembled in calmer times.
- Treat disagreement as intelligence, not failure – where two professional communities split, the disagreement is itself a finding.
- Keep AI downstream of human judgment: Use models to translate and decompose expert assessments, never to author them.
- Invest in cross-functional expert networks, not in-house teams alone: Diverse, dispersed panels outperform small, like-minded ones.
- Start building the longitudinal baseline now – a multi-year record of structured expert judgment is a time-locked asset a later entrant cannot buy.

1. The Challenge and Why Today's Tools Fall Short

The episode above was not exceptional. Since 2018, firms in globally integrated value chains have absorbed a rolling series of geopolitical shocks (e.g., trade wars, pandemic border closures, the weaponisation of energy, technology export bans, wars in Europe and the Middle East) each demanding a strategic response faster than traditional intelligence-gathering could deliver. Caldara and Iacoviello (2022) showed that spikes in geopolitical risk depress investment, trade, and employment by more than conventional models predict; Witt (2019) traced the structural reversal of globalisation that produced them. The diagnosis is well established. The operational response is not.

Two features make this environment different from the one the political-risk playbook was written for. The first is speed: a tariff announced on a Sunday can invalidate a procurement commitment by Monday; an export control published one week can strand a vendor relationship the next. Geopolitical risk now runs on a news cycle, not an investment cycle. The second is interconnection: a single US semiconductor control propagates through third-country suppliers, Southeast Asian fabrication schedules, Swiss optical-component pricing, and Dutch capital-expenditure plans at once. Events arrive not as discrete shocks but as cascades that cross industry, geographic, and functional boundaries (Chopra and Sodhi, 2004), except that these cascades now originate in government buildings rather than on factory floors.

The practical consequence is specific. The expertise needed to assess a geopolitical event's impact on a given value chain is dispersed across people in different roles, industries, and countries: the procurement director in Zurich who knows the supplier dependency, the trade analyst in Brussels who knows the regulatory trajectory, the operations manager in Shenzhen who knows the factory-level workaround, the materials scientist in Munich who knows whether substitution is feasible. No one person holds enough of the picture; no single consulting firm has all of them on staff.

"The problem is not a shortage of expertise. It is a shortage of infrastructure to connect, elicit, and aggregate it under time pressure."

The tools firms reach for today each fail in a different way. Commissioning an eight-week consulting study produces depth, but it arrives after the decision window has closed. Asking a generative-AI model for an instant assessment produces fluency, but the output has no author: it cannot be questioned in a board meeting, cited in a regulatory filing, or held accountable for being wrong. Relying on a small in-house intelligence team produces speed, but the team shares one professional formation and sits in one organisational silo, so its analysis inherits a narrow set of assumptions. What is missing is an approach that is simultaneously fast, accountable, and cognitively diverse. That is the gap the platform described here is built to close.

The rest of this paper sets out that approach. At its core is a simple division of labour, summarised in Table 1: Software detects the events that warrant attention, a structured panel of human experts judges what they mean, and AI translates that judgment into specific consequences for a specific firm. Each layer does what it does well and hands off to the next; human judgment is never displaced by machine processing. We introduce the architecture first, show it at work in two short cases, and close with five recommendations for decision-makers.

2. The Decision Intelligence Platform: A Three-Layer Architecture

“Decision intelligence” is the discipline of turning information into better decisions at scale, sitting at the intersection of data science, decision theory, and managerial practice (Pratt, 2019). The hard design question is not whether to use AI but how to allocate decision rights between human and algorithmic agents (Shrestha, Ben-Menahem and von Krogh, 2019). The platform explicitly resolves the allocation across three layers, as shown in Table 1 and used as a reference point throughout this paper.

Table 1. *The three-layer Decision Intelligence architecture*

Layer	Function	Technology	Human governance
Signal layer	Continuous environmental scanning; event detection; relevance scoring; alert triggering	Wippli Signals: A OSINT (open source intelligence) engine with relevance scoring.	Domain experts set thresholds and risk-category definitions
Assessment layer	Structured expert panels; iterative Delphi rounds; convergence tracking; disagreement profiling	NirnayaX: A Delphi study platform based on an advanced Decision Intelligence Operating System; deterministic, no (Gen)AI in core.	Researchers design frameworks, curate panels, and validate outputs
Translation layer	Business-model decomposition; scenario-to-consequence mapping; resource-impact estimation	LLM (Large Language Model) empowered workflows: User-supplied models (bring-your-own-AI) to analyze the Delphi results or pre-defined workflows developed by HSG faculty.	Strategic-management ontologies govern decomposition templates


The signal layer performs the early-warning function. Open-source intelligence (OSINT) (i.e., the systematic collection and analysis of publicly available information, from regulatory filings and trade statistics to news feeds and satellite imagery) is a discipline in its own right, not a poor substitute for classified material (Steele, 2007). The platform’s engine (e.g., Wippli Signal) continuously monitors open sources across predefined risk categories (trade policy moves, export controls, regulatory shifts, supply chain disruptions) and applies relevance scoring calibrated to specific industries and geographies. OSINT is unmatched at speed and coverage: it can flag a filing in Beijing as relevant to European semiconductor makers within minutes. What it cannot do is judge what that filing means two and three steps downstream. So, when a signal crosses its threshold, the engine does not answer the question, but it can trigger a panel that will.

The assessment layer supplies the judgment. It runs on the Delphi method, invented at RAND in the 1950s to elicit structured expert judgment on questions that could not be computed (Dalkey and Helmer, 1963). The logic is unchanged: assemble experts, collect independent judgments on a structured set of questions, aggregate them anonymously, feed the aggregate back, and invite revision. Anonymity prevents anchoring to a dominant voice; iteration forces reflection; aggregation yields a distribution of opinion rather than a single forecast, and the distribution, by showing where experts agree and where they do not, is often more informative than its average. The evidence is encouraging: Delphi panels are still perceived to outperform unstructured group discussion on most

forecasting tasks (Rowe and Wright, 1999), and while individual experts are overconfident and barely better than chance (Tetlock, 2005), the structured aggregation of independent judgments reliably beats both lone experts and prediction markets (Tetlock and Gardner, 2015). Crucially, the assessment core runs without AI: the aggregation is deterministic (i.e., not generative), and the convergence tracking is mathematical, so a board, regulator, or journalist can trace exactly how a conclusion was reached, through which framework, from how many respondents, and with what convergence. In a world increasingly sceptical of machine-generated opinion, that auditability is itself a source of credibility.

The translation layer turns judgment into action. A panel's output (e.g., a probability distribution, a convergence map, a scenario) is information or knowledge, not yet a decision. Knowing that experts rate a sustained rare-earth restriction at 70% over eighteen months tells a CEO nothing, by itself, about her procurement contracts, R&D pipeline, or capital plan. This is where large language models earn their place: not as forecasters, but as decomposition and translation engines downstream of human judgment (Agrawal, Gans and Goldfarb, 2018, 2022). The distinction matters because it matches AI to the work it does well. AI-assisted professionals do markedly better at structured analysis and writing that fall within a model's capability frontier, but worse at judgment tasks the model cannot perform, precisely because its fluency hides its errors (Dell'Acqua et al., 2023). The platform keeps the model strictly on the right side of that line: it decomposes and translates; it does not judge.

Figure 1. Integrated Decision Intelligence Platform for geopolitical risk



The integration is the innovation (Figure 1). Each component is individually mature (e.g., Delphi since the 1950s, OSINT since the 1990s, LLM text processing since 2023), but joining them lets each cover the others' blind spots. OSINT without Delphi produces alerts without judgment; Delphi without OSINT produces judgment without triggers; both without translation produce intelligence without action. Information flows upward (a signal triggers a panel; a panel produces an assessment; an assessment feeds a decomposition), while governance flows downward (humans set the alert thresholds, design

the panel frameworks, and own the decomposition templates). The machine scans, detects, and decomposes; the human judges, contextualises, and disagrees productively.

3. Two Basic Application Modes in the Context of Geopolitics: Foresight and Rapid Response

The same architecture, expert network, and database support two modes of use in the context of geopolitics, distinguished by their time horizons and triggers.

In strategic **foresight mode**, the platform addresses structural questions (e.g., how geopolitical, technological, and regulatory trends will reshape an industry or function over a three- to ten-year horizon). A typical study convenes 30+ experts across a value chain and runs two Delphi rounds over three to five weeks, producing not a single prediction but a convergence map that shows where expert opinion agrees and where it splits.

In **rapid-response mode**, the trigger is an event. When OSINT engines (e.g., Wippli Signals) detect something that crosses a relevance threshold, a tightly focused panel of 12+ experts is assembled within 24 hours and conducts an online, real-time Delphi, producing a structured brief in 48 to 72 hours (if required). Speed is possible only because the foresight infrastructure already exists: the registry of experts is built, the frameworks are designed and need only minor adaptation, rather than starting from a blank page.

The two cases below show each mode in action.

4. The Decision Intelligence Platform in Action: Two Short Cases

Case A / Foresight: critical minerals and European precision manufacturing (2027–2035)

Consider a foresight study on the concentration of critical-mineral supply chains and its implications for European precision manufacturing. The panel draws on 67 experts spanning manufacturing operations, mining, trade policy, materials science, strategy, and investment. Five structured dimensions frame the assessment: supply-concentration risk by mineral; substitution feasibility by application; regulatory and trade-policy trajectory; recycling and circular-economy potential; and competitive-repositioning options for European manufacturers.

The output is a convergence map, not a forecast. Suppose it shows that 85% of panellists expect germanium-supply diversification to accelerate, but only 40% believe it will be sufficient to eliminate European dependency within the decade. That 45-point gap between “accelerating” and “sufficient” is the finding: it tells a procurement director that diversification is real but not yet investable as a sole-source strategy — a far more actionable conclusion than either a single point estimate or a consensus that would have papered the disagreement over.

Foresight studies of this kind compound in value. Reconvene the same panel a year later and the prior convergence map becomes a baseline: did the panel’s optimism about germanium substitution survive the latest export controls? Where did opinion move, and where did it hold? Scenario planning’s real payoff was always the calibration of strategic attention — the ability to recognise, when an event

arrives, that it was anticipated and a response rehearsed (Schoemaker, 1995). A longitudinal Delphi record delivers that calibration at a precision one-off exercises cannot match.

Case B / Rapid response: from a gallium export control to a board-ready decision in five days

Now suppose that, six months after that study, a new Chinese export control on gallium compounds is published. Because the critical-minerals network already exists, the platform moves from signal to decision in days, not weeks.

Day 0. Wippli Signals detects the filing within hours of publication, scores it as high-relevance for European semiconductor and optical-component manufacturers, and triggers a pre-configured ad-hoc panel.

Days 1–2. A focused panel of 22 experts — operations managers, trade lawyers, materials scientists, and strategists drawn from the existing registry — is convened. They are not starting cold: they begin from the probability distribution the foresight study already established, so the question is narrow and answerable in twenty minutes rather than twenty hours — how does this specific control change that distribution?

Days 3–4. A single Delphi round, with a short revision window, produces a structured brief. It records not only the panel's central assessment but its disagreements — and here the split is the signal. Operations experts may rate the disruption as severe and near-term, while materials specialists rate substitution as feasible within two quarters. That divergence tells leadership exactly where the uncertainty lives and which assumption to test first.

Day 5. The translation layer converts the assessment into functional consequences. Instead of a narrative a strategy officer reads and files, each function receives a directive that traces back to the panel's judgment: procurement sees which vendor categories to re-qualify and on what timeline; R&D sees which substitution programmes to accelerate; finance sees an investment range and payback logic; the board sees a single traceable chain from “what the experts told us about the world” to “what we propose to spend, and why.”

The point of the case is not the specific numbers, which depend on the event, but the elapsed time and the accountability: five days, with a documented reasoning trail, against the eight weeks a consulting engagement would need and the unaccountable fluency an off-the-shelf model would offer.

5. From Scenario to Functional Consequence: How the Translation Layer Works

The translation in Case B is not improvised. It runs through a structured workflow (in this case, the Institutions-Resources Matrix × Business-Model-Pattern (IRM × BMP) decomposer developed at the University of St.Gallen) that converts a Delphi-validated scenario into a prioritised investment portfolio, with every output traceable back to the original judgment.

The workflow first fixes the strategic context: the scenario, the firm's identity and horizon, and two hypotheses (i.e., a default (what happens if the firm holds course) and a pivot (what the scenario might warrant)). It then derives competitive implications across five stakeholder domains: customers,

competitors, suppliers, government, and society. Next, it uses the Business Model Navigator's logic of business model patterns (Gassmann et al., 2014), identifying the two or three business-model patterns the scenario threatens, amplifies, or newly enables, because that classification sets the direction of response: defend or exit a threatened pattern, accelerate an amplified one, build a newly enabled one.

The analytical core is a six-stage decomposition of each pattern: its operating mechanism across product, data, contract, and process layers; a quantified KPI baseline; a stress test across optimistic, base, and pessimistic cases that labels each investment no-regret, conditional, or avoid; a competitive benchmark; an implementation roadmap; and a map of how the patterns interact. The final steps translate this into a Stop / Continue / Start directive for each key function and a mapping across five resource categories — physical assets, people, technologies, structures, and relationships — each with a cost, a timeline, and a priority tier.

The output is a prioritised investment portfolio, not a strategy essay. A CFO reads investment ranges and payback logic; a CHRO reads hiring and capability implications; a CTO reads technology priorities; a board member reads an unbroken chain from world to spend. Teece (2010) called the capacity to redesign a business model in response to a changing environment a dynamic capability; this workflow operationalises it.

The role of AI here illustrates the platform's governing principle. At each step the user may invoke a model — their own subscription, by design, which keeps costs predictable and avoids lock-in — to draft the competitive implications, pattern classifications, decompositions, and resource maps. But the model only fills fields the methodology defines; it never designs them, and every output is editable before it passes to the next stage. The methodology designs the fields; the AI fills them; the human decides. That separation keeps the result auditable even as the underlying models improve, and it guards against the well-documented tendency of decision-makers to stop at an appealing scenario and never test it (Kahneman, Lovallo and Sibony, 2011).

6. Key Takeaways for Executives

Five practical implications follow for any organisation seeking to build the capability to navigate geopolitical disruption.

1. **Build intelligence infrastructure before you need it.** The five-day response in Case B is possible only because the expert registry, analysis frameworks, and related components were assembled in advance. A firm that waits for the crisis to build them will take five weeks - by which time the window has closed. The infrastructure is the capability.
2. **Treat disagreement as intelligence, not failure.** The instinct to converge on a single recommended view destroys information. Where two professional communities split on a scenario's feasibility, the gap reveals an assumption mismatch that consensus-seeking would have hidden - and that is often the most decision-relevant thing in the brief.

3. **Keep AI downstream of human judgment.** Use models to translate and decompose, never to author the assessment. An LLM-generated risk call has no author: it cannot be questioned, cited, or held accountable. The expert panel supplies the provenance a board and a regulator require; the model supplies the speed of decomposition.
4. **Invest in cross-functional expert networks, not in-house teams alone.** A panel drawn from an alumni network, an industry consortium, or a cross-functional registry aggregates judgments from people in different roles, sectors, and countries, breaking the shared assumptions that constrain a small internal team. Panel composition is a strategic choice - and it matters more than the sophistication of the tool.
5. **Start building the longitudinal baseline now.** After three years, the platform's most valuable asset is neither the software nor the registry but the longitudinal database of structured expert judgment. It is a time-locked asset: a competitor starting in 2029 cannot buy the 2027–2029 baseline. Beginning now is the decision that compounds.

7. Conclusion: Decision Intelligence as Institutional Capability

The argument compresses into a sentence: the expertise to navigate geopolitical disruption already exists, but the infrastructure to mobilise, aggregate, and translate it does not, and building that infrastructure is now an urgent priority for firms, universities, and the institutions that connect them.

The platform described here is one possible design, not the only one, and it has limits worth stating plainly. Delphi panels are only as good as the experts in them, and expert availability is a recurring constraint. OSINT triggers need careful calibration to avoid both alert fatigue and missed signals. The translation tools are powerful but require human validation before they inform real resource decisions. None of these limits is fatal; none should be minimised.

What justifies the effort is the alternative. The status quo (e.g., phoning three contacts, commissioning studies that arrive too late, or trusting fluent and unaccountable machine summaries) is not free. It is paid for in misjudged commitments, missed regulatory transitions, delayed pivots, and the slow erosion of position that follows from underestimating political risk (Henisz and Zelner, 2010).

The deeper question is institutional: who should own such infrastructure? A single firm draws on a too narrow expert base; a consulting firm faces an incentive to tell clients what they want to hear; a university brings independence but risks being too slow. A cross-functional, cross-geographic network (e.g., an alumni body) may offer the most promising foundation if it can build the governance, technology, and culture of structured disagreement that the task demands.

And that points to what a research university is for in an age when information is abundant and judgment is scarce. If the university's distinctive contribution is no longer the production of information (which AI now produces faster and cheaper) but the cultivation of judgment, the discipline of structured disagreement, and the independence to report what the evidence says rather than what the audience wants to hear, then a decision intelligence platform is not a side project. It is a

university's core mission, rendered as operational infrastructure that the world outside urgently needs.

References

- Agrawal, A., Gans, J. and Goldfarb, A. (2018) *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Agrawal, A., Gans, J. and Goldfarb, A. (2022) *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Caldara, D. and Iacoviello, M. (2022) 'Measuring Geopolitical Risk,' *American Economic Review*, 112(4), pp. 1194–1225.
- Chopra, S. and Sodhi, M.S. (2004) 'Managing Risk to Avoid Supply-Chain Breakdown,' *MIT Sloan Management Review*, 46(1), pp. 53–61.
- Dalkey, N. and Helmer, O. (1963) 'An Experimental Application of the Delphi Method to the Use of Experts,' *Management Science*, 9(3), pp. 458–467.
- Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, R. (2023) 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality.' *Harvard Business School Working Paper*, No. 24-013.
- Gassmann, O., Frankenberger, K. and Csik, M. (2014) *'Business Model Navigator'*. Pearson Education Limited.
- Henisz, W.J. and Zelner, B.A. (2010) 'The Hidden Risks in Emerging Markets,' *Harvard Business Review*, 88(4), pp. 88–95.
- Kahneman, D., Lovallo, D. and Sibony, O. (2011) 'Before You Make That Big Decision,' *Harvard Business Review*, 89(6), pp. 50–60.
- Pratt, L. (2019) 'Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better World.' Emerald Publishing.
- Rowe, G. and Wright, G. (1999) 'The Delphi Technique as a Forecasting Tool: Issues and Analysis,' *International Journal of Forecasting*, 15(4), pp. 353–375.
- Schoemaker, P.J.H. (1995) 'Scenario Planning: A Tool for Strategic Thinking,' *Sloan Management Review*, 36(2), pp. 25–40.
- Shrestha, Y.R., Ben-Menahem, S.M. and von Krogh, G. (2019) 'Organizational Decision-Making Structures in the Age of Artificial Intelligence,' *California Management Review*, 61(4), pp. 66–83.
- Steele, R.D. (2007) 'Open Source Intelligence,' in Goldman, J. (ed.) *Ethics of Spying: A Reader for the Intelligence Professional*. Lanham, MD: Scarecrow Press, pp. 148–164.
- Teece, D.J. (2010) 'Business Models, Business Strategy and Innovation,' *Long Range Planning*, 43(2–3), pp. 172–194.
- Tetlock, P.E. (2005) *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton: Princeton University Press.
- Tetlock, P.E. and Gardner, D. (2015) *Superforecasting: The Art and Science of Prediction*. New York: Crown.
- Witt, M.A. (2019) 'De-globalization: Theories, Predictions, and Opportunities for International Business Research,' *Journal of International Business Studies*, 50(7), pp. 1053–1077.