

# How Foresight Makes a Country Grow

Foresight is not a report — it is a standing capability: built as an engine, run continuously by AI agents, it converts the cost of being surprised into the option value of acting early.

# 32 areas

ONE FORESIGHT CAPABILITY, APPLIED ACROSS 32 NATIONAL DOMAINS AND RANKED IN FOUR TIERS OF LEVERAGE · ENSI FORESIGHT DIVISION, 188-SOURCE LIBRARY

# ≈3×

THE LONG-RUN COST OF REACTING LATE VS. ACTING EARLY, COMPOUNDED ACROSS DOMAINS — THE CENTRAL WAGER OF THIS REPORT

ISSUE 2026    188 PRIMARY SOURCES    20 RESEARCH ANGLES    32 AREAS    4 TIERS

5 / 15 / 30-YR HORIZONS

STRATEGY REPORT Nº03 · THE FORESIGHT DIVISION

# One engine, four tiers, thirty-two applications — and how to read it at four depths.

This is an ENSI Foresight Division analysis, the third in a sequence. Report Nº01 made the growth case; Report Nº02 made the agentic-operations case. This report fuses them: it argues that a country's single highest-return investment is a **standing foresight capability** — and that AI agents are what finally make such a capability continuous, affordable, and whole-of-government in scope.

### The engine

**Scan → model → stress-test → route-to-decision.**  
 A continuous loop that reads the global signal space, builds and refreshes scenarios, attacks every plan, and pushes the relevant slice into the decision being made anyway — the budget, the legislative programme, the security review.

### The discipline

**Human-in-the-loop, audit, provenance.** Agents widen the aperture; humans keep the accountability. Every signal is traceable to source, every judgement is owned by a named officer, and every model run is logged — so the capability earns trust rather than assuming it.

### THE REFRAME THAT ORGANISES EVERYTHING

Most governments are run as if the future were a rumor — budgeted one year at a time, surprised expensively by forces that were knowable all along. Foresight is not the art of predicting the unpredictable. It is the discipline of refusing to be surprised by the predictable, and staying nimble in the face of the genuinely uncertain. That is a **capability**, not a document.

### FOUR READING PATHS

DEPTH	TIME	WHAT YOU GET	WHERE
Depth 1	90 seconds	The whole thesis — cover, executive summary, the dashboard, the engine in one diagram, twelve findings.	§1 · §4–§7
Depth 2	20 minutes	The nine analytical angles — the report's main contribution, where the understanding actually forms.	§8–§16
Depth 3	reference	The 32-area catalog — each a seven-part operating brief, grouped into four tiers of leverage.	§17–§20
Depth 4	the proof	Cases, the 12-month build roadmap, the steelman of objections, methodology, the 188-source library.	§21–§26

**03**

ENSI Strategy Report in the growth → agents → foresight sequence

**33**

sections across six movements

**≈40%**

of the report is original analysis (Movement B)

**5**

agent archetypes run the engine

## CONTENTS

# The architecture of the argument.

Front-loaded by design: a reader who never reaches the catalog still owns the full thesis. Movements A and B carry the intellectual payload; **Movement F** then specifies the engine itself — the Agentic Foresight Architecture, 26 components in four layers.

## A · THE 90-SECOND READ

01	Cover — the defining claim	1
02	Masthead & how to read	2
03	Contents	3
04	Executive summary	4
05	The Foresight Dashboard — thesis in numbers	6
06	The Foresight Engine — one diagram	8
07	Twelve findings	9

## B · THE ANALYSIS — NINE ANGLES

08	Reframe: the future is not a rumor	10
09	The Foresight Capability explained	12
10	The economics of anticipation	14
11	Cross-cutting patterns	16
12	The minimum viable foresight state	19
13	The agentic shift & guardrails	21
14	The four tiers mapped	24
15	How the best foresight states do it	26
16	Why foresight units die	29

## C · THE 32 AREAS (REFERENCE)

17	Tier I — The Foresight Core (1–8)	31
----	-----------------------------------	----

18	Tier II — Productive Capacity (9–16)	36
19	Tier III — Society & Quality of Life (17–24)	40
20	Tier IV — Environment, Security & Enablers (25–32)	44

## D · APPLICATION

21	Cases: foresight states working	48
22	The 12-month build roadmap	50
23	Risks, counter-arguments & steelman	52

## F · THE AGENTIC FORESIGHT ARCHITECTURE

24	The engine underneath — the forecastable curve	58
25	The 26 components at a glance	60
26	The seven agent archetypes	61
27	The four layers — Core · Data · Fleets · Governance	62
28	The living substrate — how the loop runs	66
29	The 102-source technology-foresight library	67
30	First 24 months — the build sequence	68

## E · BACK MATTER

31	Methodology	69
32	The evidence base: 188-source library	70
33	Glossary, references & back cover	72

### THE SPINE — REPEATED AT EVERY DEPTH

A country does not need better forecasts; it needs a standing **foresight capability** — scan → model → stress-test → route-to-decision, run continuously by AI agents — because the cost of acting early is a fraction of the cost of reacting late, and that gap, compounded across 32 domains, is the difference between a country carried by the future and one that authors it.

EXECUTIVE SUMMARY · 1 OF 2

# The future was knowable; the surprise was a choice.

Most governments are run as if the future were a rumor. Budgets are set one year at a time, ministers are rewarded for what lands inside an electoral cycle, and the slow, certain forces that actually decide a country's fate — ageing, the energy transition, the arrival of machine intelligence, the drift of supply chains — are treated as someone else's problem until they arrive as a crisis. Then the same governments discover, expensively, that the future was knowable all along. The pandemic was on every risk register. The energy shock was in every scenario. The skills gap was a demographic certainty twenty years out.

The reframe that organises this whole report is simple and, in practice, radical: **foresight is not a report, it is a capability** — a standing function that continuously scans, models, and stress-tests the country's possible futures, and routes what it learns into the decisions being made anyway. The states that have built this — Singapore, Finland, the UAE, Canada — did not get better forecasts than everyone else. Forecasts are usually wrong. What they got was **option value**: the ability to see a shift early, to have already rehearsed the response, and to act while action is still cheap.

Why a country, and not a company? Because the state is the only actor with the time horizon, the breadth, and the obligation to think in decades. A firm that misreads the future loses money and exits; a country that misreads the future loses a generation. And foresight's payoff is highest exactly where markets are blindest — in long-lived infrastructure, in public goods, in systemic risk, in the slow accumulation of national capabilities. That is why a foresight capability is a growth strategy, a stability strategy, and a quality-of-life strategy at once: it lets a country invest ahead of demand, absorb shocks rather than be broken by them, and choose the future it wants rather than inherit the one it failed to see coming.



*Foresight is not the art of predicting the unpredictable; it is the discipline of refusing to be surprised by the predictable — and of staying nimble in the face of the genuinely uncertain.*

— ENSI FORESIGHT DIVISION, **THE GOVERNING REFRAME**

**32**

national domains, one engine applied 32 ways

**4**

tiers, ranked by downstream leverage

**5/15/30**

the three decision horizons foresight informs

**188**

primary sources behind the analysis

EXECUTIVE SUMMARY · 2 OF 2

# What makes this an engine, and what makes it run.

What is new — and what makes this an ENSI document rather than a restatement of classical futures studies — is the **agentic engine**. Traditional foresight is bottlenecked by human attention: a handful of analysts can scan only so many signals and refresh them only so often, so foresight arrives as an occasional, expensive set-piece.

AI agents dissolve that bottleneck. A standing fleet of **scanning agents** can read the entire global information space every day; **scenario agents** spin and stress-test hundreds of futures; **simulation agents** run agent-based models of the economy, the labour market, or an epidemic; **red-team agents** attack every plan; and **briefing agents** deliver the relevant slice to the right decision-maker at the moment of decision. Foresight stops being a quarterly artefact and becomes a continuous, living nervous system for the state — with humans keeping every judgement and every accountability.

The 32 areas are **ranked by leverage** and grouped in four tiers. Tier I is the foresight core — the central capability plus the macro-fiscal, growth, technology, labour, demographic, energy and risk functions every other decision hangs from. Tier II is productive capacity and competitiveness. Tier III is society, resilience and quality of life. Tier IV is environment, security, and the enabling layer — the data architecture and the futures-literacy the whole system runs on. The ranking is not just a list; it is a **sequencing instruction**: build the engine once, on a shared data backbone, staffed by people who can think in futures, then point it in turn at each tier.

WHAT TO KNOW IN FIVE POINTS

- 1 **The default is failure.** Reactive government discovers the future as a crisis, when options are worst and costs highest.
- 2 **Foresight is a capability, not a report.** A standing function that scans, models, stress-tests, and routes — continuously, for the whole of government.
- 3 **Agents make it continuous.** Five agent archetypes dissolve the human-attention bottleneck; humans keep judgement, audit and provenance.
- 4 **Analysis comes first.** The payoff is option value and institutional reflexes, not better point forecasts — and it is forecastable where markets are blind.
- 5 **It is growth, stability and quality of life at once.** One engine, 32 domains; the order you build it in is the strategy.

## 1.8.31.32

The minimum viable foresight state — core, risk, data backbone, futures-literacy — built first

ENSI sequencing

## 5 archetypes

scan · scenario · simulate · red-team · brief — the agents that run the engine

ENSI framework

## 3 layers

decisions informed at 3–5 yr, 10–15 yr and 25–30 yr horizons

ENSI horizons

THE FORESIGHT DASHBOARD · 1 OF 2

# The thesis, in numbers.

Twelve figures and two charts that carry the argument before a word of analysis. Facts we own are tagged ENSI; institutional figures carry their source.

<p><b>32</b> areas</p> <p>national domains where anticipation changes the outcome</p> <p>ENSI catalog</p>	<p><b>4</b> tiers</p> <p>ranked by downstream leverage, core → enablers</p> <p>ENSI ranking</p>	<p><b>188</b></p> <p>primary sources — foresight toolkits, the data/statistics literature, the economics of growth</p> <p>ENSI library</p>	<p><b>20</b></p> <p>distinct research angles synthesised into the analysis</p> <p>ENSI method</p>
<p><b>≈3×</b></p> <p>indicative ratio of react-late to act-early cost, compounded across domains — the core wager</p> <p>ENSI thesis</p>	<p><b>7</b> parts</p> <p>the operating brief behind every area: short · why · questions · signals · methods · agents · wiring</p> <p>ENSI catalog</p>	<p><b>5</b></p> <p>agent archetypes that run the continuous engine</p> <p>ENSI framework</p>	<p><b>3</b></p> <p>decision horizons informed — 5, 15 and 30 years</p> <p>ENSI horizons</p>
<p><b>~\$3.8<sup>tn</sup></b></p> <p>estimated global output lost to the 2020 pandemic shock through 2021 — a risk that sat on every register</p> <p>IMF / World Bank est.</p>	<p><b>8–12</b></p> <p>staff in a lean central foresight directorate — the whole engine, stood up in year one</p> <p>ENSI · Area 1</p>	<p><b>~2<sup>pp</sup></b></p> <p>order-of-magnitude GDP swing foresight contests across the macro spine over a cycle</p> <p>ENSI estimate</p>	<p><b>2-pg</b></p> <p>the weekly anticipatory brief — agents draft, humans own — that routes foresight into power</p> <p>ENSI · Area 1</p>

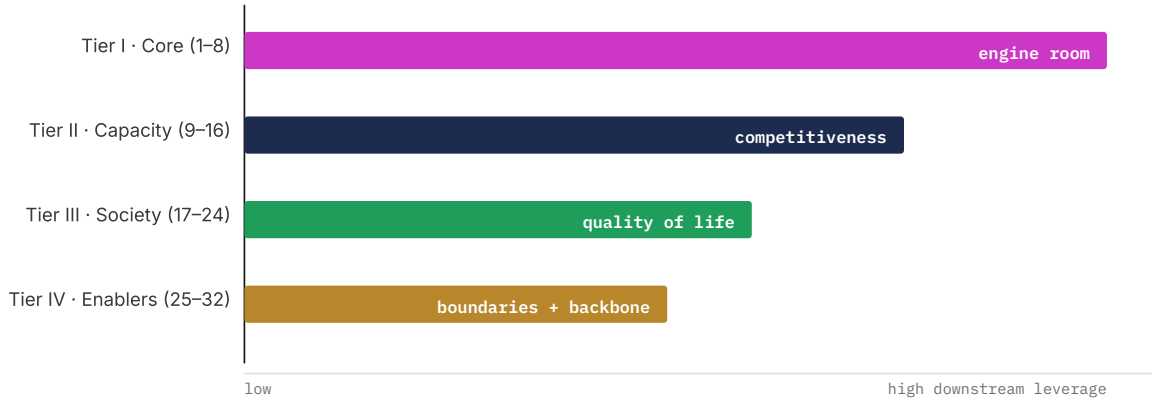
READ THE TILES AS ONE CLAIM

A modest, standing cost — a dozen people and an agent fabric — set against the large, lumpy cost of being surprised: a skills gap that becomes mass unemployment, a risk that becomes a crisis, an asset that becomes a stranded liability. That asymmetry, not any single number, is the thesis.

THE FORESIGHT DASHBOARD · 2 OF 2


# Where leverage sits, and why early is cheap.

FIG. 1



**Fig. 1 — Leverage ranking.** Tiers ordered by how much getting foresight right changes everything downstream. The core ranks first not because it matters more to citizens, but because every other tier inherits its advantage — or its absence. Bar lengths are illustrative of rank, not measured magnitudes.

FIG. 2

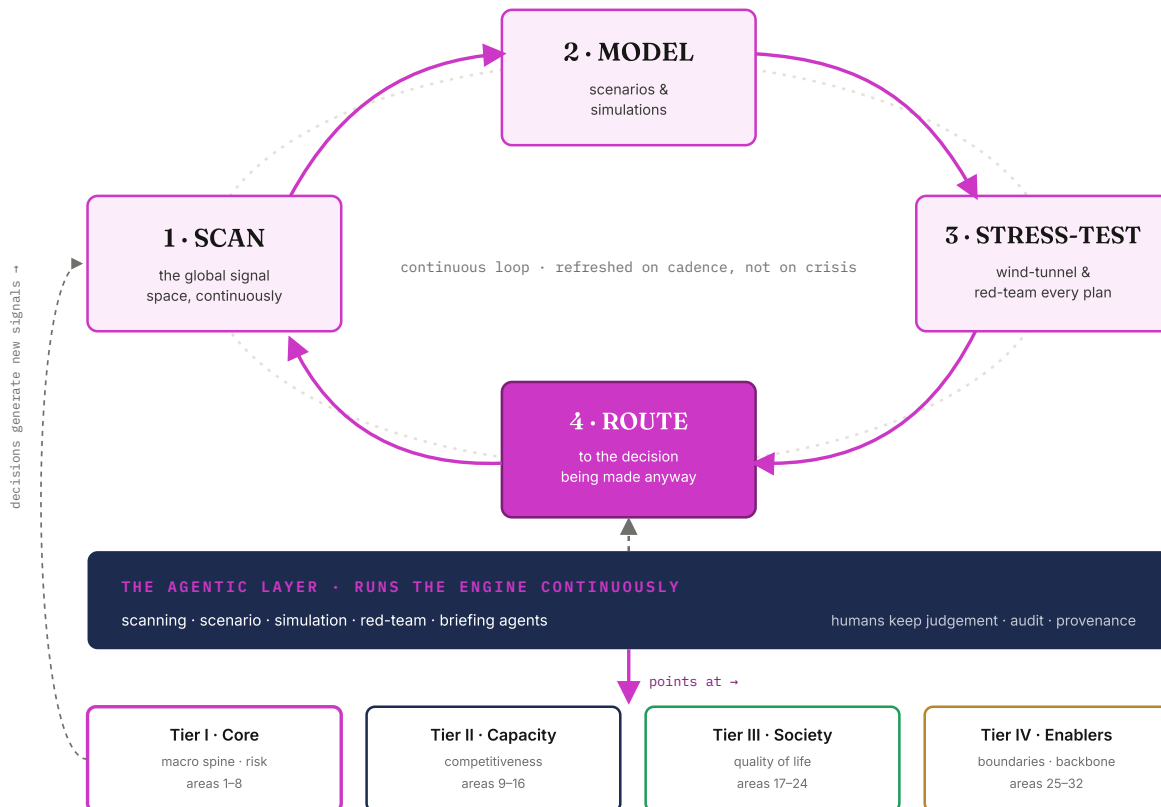


**Fig. 2 — Act early vs react late.** The cost of acting on a shift — a technology, a skills gap, a fiscal liability — rises steeply along its diffusion curve. Intervene inside the green window, while options are cheap and reversible; intervene late and you buy influence over an entrenched system at ruinous cost. The magenta gap, compounded across 32 domains, is the entire economic case for foresight. Schematic.

THE MODEL IN ONE PICTURE

# The Foresight Engine.

One continuous loop — scan → model → stress-test → route-to-decision — run by an agentic layer and pointed, in turn, at the four tiers of national life. Build the engine once; apply it 32 ways.



**Fig. 3 — The Foresight Engine.** The same four-step loop runs in every one of the 32 areas; only the domain changes. An agentic layer of five archetypes runs it continuously, with humans owning judgement, audit and provenance. The engine is built once — at the centre, on a shared data backbone — then pointed in turn at the four tiers, whose decisions feed fresh signals back into the scan. That is the practical meaning of foresight as a capability rather than a report.

## THE 90-SECOND READ · SCANNABLE PROOF

# Twelve findings.

The argument compressed to twelve one-sentence claims, each carrying the evidence chip that backs it.

- 1 **Reactive government is the default failure mode.** Budgets set one year at a time discover the future as a crisis, when options are worst and borrowing costs highest. [Reframe · §8](#)
- 2 **Foresight is a capability, not a report.** The states that lead — Singapore, Finland, the UAE, Canada — did not get better forecasts; they built a standing function that routes anticipation into decisions. [Benchmark · §15](#)
- 3 **The payoff is option value, not accuracy.** Forecasts are usually wrong; what foresight buys is the ability to see a shift early, rehearse the response, and act while action is still cheap. [Economics · §10](#)
- 4 **The asymmetry is the whole case.** A standing cost of a dozen people and an agent fabric is set against the lumpy cost of being surprised — measured in points of GDP and avoided crises. [≈3× react-late vs act-early](#)
- 5 **One engine runs all 32 areas.** Scan → model → stress-test → route is identical across domains; only the signals and the owner change. [Framework · §9](#)
- 6 **Agents dissolve the attention bottleneck.** Five archetypes — scan, scenario, simulate, red-team, brief — turn foresight from a quarterly set-piece into a continuous nervous system. [5 archetypes · §13](#)
- 7 **Discipline is what earns trust.** Humans keep every judgement, with audit and provenance on every signal and run; agents widen the aperture, never the accountability. [Guardrails · §13](#)
- 8 **The ranking is a sequencing instruction.** Areas 1, 8, 31 and 32 — core, risk, data backbone, futures-literacy — are the minimum viable foresight state; build them first. [1·8·31·32 · §12](#)
- 9 **The state is the right actor.** Only government has the horizon, breadth and obligation to think in decades — and foresight pays most where markets are blindest. [Patterns · §11](#)
- 10 **It is a growth and stability and quality-of-life story at once.** One capability lets a country invest ahead of demand, absorb shocks, and choose its future. [Thesis · §4](#)
- 11 **The expensive surprises are predictable.** The pandemic was on every register, the energy shock in every scenario, the skills gap a demographic certainty — foresight refuses to be surprised by the knowable. [188 sources](#)
- 12 **Foresight units die for structural reasons.** Orphaned reports, no routing, no data, no mandate — the failure taxonomy is known, which means it is avoidable. [Failure modes · §16](#)

**THE LINE THAT HOLDS IT TOGETHER**

The wager of this report is that anticipation is the cheapest investment a state can make — and that the gap between acting early and reacting late, compounded across 32 domains, is the difference between a country carried by the future and one that authors it.

REFRAME · 1 OF 9 ANGLES

# Government runs as if the future were a rumor — and that is the default, expensive failure.

Budgets are set one year at a time. Ministers are rewarded for what shows up inside an electoral cycle. The long, slow, certain forces that actually decide a country's fate are treated as someone else's problem — until they arrive as a crisis, and the state discovers, expensively, that the future was knowable all along.

The pandemic was on every risk register. The energy shock was in every scenario set. The skills gap was a demographic certainty visible twenty years out. None of these were unknowable; all of them were unattended. The failure was not of forecasting but of **posture** — a state organised to react to the present cannot see the future even when the future is plainly written down. Ageing, the energy transition, the arrival of machine intelligence, the drift of supply chains: each moves on a clock measured in decades, and each is routinely ignored on a clock measured in months.

This is not a failure of individual ministers. It is a structural property of how a government allocates attention. The annual budget rewards the extrapolation of the year that just ended. The legislative programme rewards what can be delivered before the next election. The career incentive rewards the absence of visible failure on your watch, not the avoidance of a crisis that lands on your successor's. The predictable, slow-moving threat is precisely the one no part of the machine is built to act on while action is still cheap — and so the cheap window closes, unnoticed, every single time.



*Foresight is not the art of predicting the unpredictable. It is the discipline of refusing to be surprised by the predictable — and of staying nimble in the face of the genuinely uncertain.*

— ENSI FORESIGHT DIVISION · THE ORGANISING REFRAME

That sentence does most of the work in this report. It splits the future into two halves and assigns each a different obligation. The **predictable** half — demography, asset lifespans, diffusion curves, fiscal arithmetic — carries an obligation to *provision early*, because the cost of acting in time is a fraction of the cost of reacting late. The **genuinely uncertain** half — shocks, tipping points, technological discontinuities — carries a different obligation: to stay rehearsed and nimble, so that when one branch resolves, the response is already on the shelf. A state that confuses the two — demanding certainty before it acts on the predictable, or freezing before the uncertain — fails in both directions at once.

## 1 yr

The horizon most budgets are actually set against

## 5–30 yr

The horizons that decide growth, stability and quality of life

## 20 yr

Lead time on a demographic certainty — known, unattended

## days

The lead time of the shock that finally forces the decision

## The myth, and the mechanism that actually operates underneath it

Every comfortable belief a government holds about the future is paired with a mechanism that quietly contradicts it. The myths are reassuring because they license inaction; the mechanisms are inconvenient because they demand provisioning ahead of demand. Naming both, side by side, is the first act of foresight — because you cannot stress-test a plan whose hidden assumption you have never said out loud.

THE COMFORTABLE MYTH	THE MECHANISM THAT ACTUALLY OPERATES
"The future is unknowable, so planning for it is guesswork."	Most of what matters is not unknowable — it is unattended. Ageing, asset lifespans and diffusion curves are forecastable to a tight range; the failure is acting on them, not knowing them.
"We'll deal with it when it arrives."	By arrival the cheap options have expired. A second supplier, a trained cohort, a retrofitted grid take years; the shock takes days. Reaction buys the worst options at the highest price.
"Forecasts are usually wrong, so foresight is a waste."	True — and beside the point. The payoff is not an accurate forecast but <b>option value</b> : seeing the shift early, having rehearsed the response, acting while action is still cheap.
"Foresight is a report we commission every few years."	A report ages the day it ships and routes nowhere. The thing that works is a <b>standing capability</b> — continuous scanning, modelling and routing into decisions being made anyway.
"The risk register covers us."	A static binder reviewed annually is contradicted by the next surprise. Real crises cascade and compound; what's needed is a live system that catches interactions a register lists separately.
"Strategic sectors are obvious; we know our strengths."	Comparative advantage migrates. Yesterday's strength is often nostalgia dressed as strategy; the reachable high-value rung is defined by where demand and technology will be in fifteen years.
"This is for big states with slack to spare."	The asymmetry is sharpest for a mid-sized state with little fiscal slack: you cannot buy your way out of a crisis you failed to imagine. Anticipation is the affordable option, not the luxury one.

### THE REFRAME, STATED ONCE

Foresight is a **capability, not a report**. The deliverable is not a document predicting the future but a standing function that continuously asks, before money is committed, "and if the world turns out otherwise?" — and routes the answer into the budget, the legislative programme, the procurement, the strategy.

### THE COST OF THE DEFAULT

The expensive surprises are almost never unknowable. They are the contingent liability nobody priced, the chokepoint nobody mapped, the cohort nobody trained — each a predictable force met with a posture built only to react. The bill is paid in points of GDP and lost years.

The rest of this analysis takes the reframe at its word. If foresight is a capability rather than a report, then the next question is mechanical: **what is the capability made of?** Not which crystal ball, but which engine — what does it ingest, how does it process, and where does its output go so that it changes a decision rather than decorating a shelf. That is the subject of the angle that follows.


FRAMEWORK · 2 OF 9 ANGLES

# Every effective foresight system is the same engine, pointed at a different domain.

Singapore's risk-and-horizon machinery, Finland's megatrend work, Canada's disruption scanning — they look different on the surface and run the identical loop underneath: **scan** → **model** → **stress-test** → **route-to-decision**, refreshed continuously rather than commissioned occasionally. Learn the engine once and you can aim it at all thirty-two domains.

FIG. B2 · THE FORESIGHT ENGINE – ONE LOOP, RUN CONTINUOUSLY

THE AGENTIC LAYER runs the loop continuously – scanning · scenario · simulation · red-team · briefing agents



**One engine, four stages, a closing loop.** Traditional foresight runs this loop once every few years by hand. The agentic layer runs it every day — which is the difference between a quarterly artifact and a living nervous system for the state.

## The four stages

- 1 **Scan.** Maintain one curated, deduplicated national signals registry — trend-breaks, weak signals, frontier-firm behaviour — drawn once for the whole of government instead of thirty private spreadsheets.
- 2 **Model.** Recombine signals into a small set of canonical scenarios and quantitative models — agent-based simulations, complexity maps, fan charts — that make the futures concrete enough to test against.
- 3 **Stress-test.** Wind-tunnel every major plan against the full scenario set; red-team it for the assumption that, if false, breaks it; pre-mortem the crisis before it happens.
- 4 **Route.** Deliver the relevant slice to the right decision-maker at the moment of decision — into the budget, the law, the procurement — so foresight is a required input, not a courtesy.

## Why the loop, not the line

A report is a line: it ends. A capability is a loop: each decision and each new signal feed the next scan, so the picture is never stale and the response is always rehearsed. The loop is what converts forecasting — which is usually wrong — into **option value**, which compounds.

**THE WHOLE FRAMEWORK IN ONE LINE**  
 One engine — **scan** → **model** → **stress-test** → **route** — run continuously by agents, pointed at four tiers and thirty-two domains. Everything in this report is an application of that single sentence.

## Where the engine is pointed: four tiers, by leverage

The same engine produces wildly different value depending on where it is aimed. The thirty-two applications are ordered by **leverage** — how much getting foresight right in that domain changes everything downstream — and grouped into four tiers. The order is not decoration; it is the build sequence. Get Tier I right and every tier above it inherits the advantage; get it wrong and no downstream foresight can compensate.

TIER	WHAT IT GOVERNS	THE CONTRIBUTION, AND WHY IT SITS HERE
I	<b>The Foresight Core</b> Areas 1–8	The central capability plus the macro-fiscal, growth, technology, labour, demographic, energy and risk functions — the spine every other decision hangs from. The engine room.
II	<b>Productive Capacity</b> Areas 9–16	Innovation, trade, digital, financial, infrastructure, materials, the firm base and investment — whether the economy keeps climbing the complexity ladder.
III	<b>Society &amp; Quality of Life</b> Areas 17–24	Health, education, ageing, migration, cohesion, regional balance, housing and wellbeing — where foresight most directly touches how people actually live.
IV	<b>Environment, Security &amp; Enablers</b> Areas 25–32	Climate, food and water, nature, defence, cyber, regulation — and, load-bearing for all the rest, the data architecture and the human futures-literacy the whole system runs on.

## The agentic layer — what makes the loop continuous

Classical foresight is bottlenecked by human attention: a handful of analysts can scan only so many signals, build only so many scenarios, and refresh them only so often, so foresight arrives as an occasional, expensive set-piece. A standing fleet of agents dissolves that bottleneck — and turns the same loop into a daily organ of the state. Five archetypes run the engine; humans curate, judge and own what goes up.

### The five agent archetypes

**Scanning agents** read the global information space every day and flag trend-breaks the moment they cross a threshold. **Scenario-generation agents** spin and refresh hundreds of futures and pressure-test the canonical set for staleness. **Simulation / ABM agents** run the economy, labour market or an epidemic across synthetic populations. **Red-team agents** adversarially attack every plan, hunting the assumption that breaks it. **Briefing & early-warning agents** compress it all into the decision-shaped brief and fire when a leading indicator crosses a line.

32

Applications of one engine

4

Tiers, ordered by leverage

### AGENTS WIDEN THE APERTURE; HUMANS KEEP ACCOUNTABILITY

The fleet reads more, models more, and refreshes more often than any human team can — but the foresight directorate curates the signal, judges significance, and owns every decision that reaches the prime minister. Capacity scales; responsibility does not move.

ECONOMICS · 3 OF 9 ANGLES

# Anticipation is the cheapest investment a state can make.


A foresight capability is a modest standing cost — a few dozen people and a fleet of agents — set against the large, lumpy, unbudgeted cost of reacting late. The entire economic case is one asymmetry: the cost of acting early is almost always a fraction of the cost of reacting late, and that gap, compounded across thirty-two domains, is measured in points of GDP.

The mistake most governments make is to treat foresight as a cost to be justified against an uncertain benefit. It is the opposite. Foresight is the purchase of **option value** — the right, but not the obligation, to act early on a future that may or may not arrive. Like any option, it is cheap relative to the loss it insures against, and its value is highest exactly where uncertainty is highest and the underlying stakes are largest. A country that holds the option can move while moving is still cheap; a country that does not must buy the underlying outright, at crisis prices, on the worst possible day.

The arithmetic is an arithmetic of asymmetries. In each case the early cost is small, recurring and chosen; the late cost is large, lumpy and forced. Three concrete shapes recur across the catalog:

THE PREDICTABLE SHIFT	ACTED ON EARLY (CHEAP, CHOSEN)	REACTED TO LATE (LUMPY, FORCED)
<b>A skills gap</b> visible a decade out in the demographic and exposure data	Redirect training places and apprenticeship pipelines toward forecast demand — a budget reallocation, not new money	Becomes <b>structural mass unemployment</b> — transfer payments, lost output, a region in decline, a political emergency
<b>A risk</b> sitting quietly on the register with a faint early signal	A day of warning and a rehearsed response — preparedness at modest standing cost	Becomes a <b>national crisis</b> — emergency spending, cascading failure, the catastrophic cost of surprise
<b>An asset</b> built for a world that is ending — a chokepoint, a fuel, a plant	Stage the bet, qualify a second source, keep the option to switch — real-options discipline	Becomes a <b>stranded asset</b> — capital written off, a dependency weaponised, a decade of lock-in

FIG. B3 · ACT EARLY VS REACT LATE – COST OVER THE DIFFUSION CURVE



**The shape is the argument.** The cost of acting on a shift rises steeply over its diffusion curve; the cost of standing watch stays low and flat. The magenta gap is what foresight buys — and it widens the longer the decision is deferred.

## Why the asymmetry is sharpest for a mid-sized state

The instinct is that anticipation is a luxury for large, rich states with slack to spare. The reverse is true. The asymmetry between early and late cost is **sharpest** precisely where fiscal slack is thinnest, because a state that cannot afford to buy its way out of a crisis cannot afford to be surprised by one. For an open, mid-sized economy embedded in larger value chains, a distant chokepoint becomes a domestic factory standstill within weeks, and a demographic spending wave that was visible for thirty years arrives with the budget already committed elsewhere. Anticipation is not the expensive option here. It is the only affordable one.

<p><b>3–10<sub>yr</sub></b></p> <p>Lead time to build a competent technician or a qualified second supplier — far longer than any electoral cycle, which is why early action is the only cheap action.</p> <p>ENSI · Areas 5 &amp; 10</p>	<p><b>20–40<sub>yr</sub></b></p> <p>Working life of energy and grid assets. A wrong bet commissioned today locks in cost or stranded-asset risk for decades — the case for staging over committing.</p> <p>ENSI · Area 7</p>	<p><b>1<sub>day</sub></b></p> <p>In a crisis, the value of a single day of early warning is enormous — and it decays fast, while the standing cost of preparedness is modest and the cost of surprise catastrophic.</p> <p>ENSI · Area 8</p>
---	--	--

Notice what the option-value frame does to the usual debate. The objection to foresight is always "but forecasts are wrong" — and the objection is correct and irrelevant. You do not buy an option because you are certain the future arrives; you buy it because you are **not**. The value lives in the asymmetry of payoffs, not the accuracy of the prediction. A foresight capability that is wrong about which scenario lands, but right that *some* scenario will, still pays for itself many times over — because it kept the response cheap, staged and reversible while everyone else was committing.

### The standing cost is small and known

A foresight capability is a lean directorate of eight to twelve, a signals registry, and a fleet of agents that does the high-volume scanning and modelling no human team could. It is a recurring line item measured in the low millions — set against contingent liabilities, stranded assets and crisis spending measured in points of GDP. The cost side of the ledger is the easy side.

### The avoided cost is large and lumpy

The expensive surprises are never the headline number. They are the off-balance-sheet liability nobody priced, the chokepoint nobody mapped, the cohort nobody trained — each a predictable force met with a posture built only to react. Foresight moves those costs from "forced and large" to "chosen and small."



*You cannot afford to buy your way out of a crisis you failed to imagine. Imagining it early is the cheapest line in the budget.*

– ENSI FORESIGHT DIVISION

#### THE ONE-LINE VERSION

Foresight converts the **cost of being surprised** — large, lumpy, forced, paid at crisis prices — into the **option value of acting early** — small, recurring, chosen, paid while action is still cheap. That conversion, compounded across thirty-two domains, is the difference between a country carried by the future and one that authors it.

ANGLE 04 · CROSS-CUTTING PATTERNS

# It is not 32 projects. It is one capability, applied 32 ways.

Read the catalog end to end and a reader expects thirty-two separate programmes — a fiscal model here, an energy dashboard there, a skills observatory somewhere else. That reading is wrong, and the mistake is expensive. Beneath the thirty-two briefs runs a single machine. Only the domain changes; the engine, the backbone, the agents and the horizons are the same. Build it once and every new area is a low-cost extension of a system that already works.

The thirty-two areas were written to be read independently, each a complete seven-part operating brief. But they were **designed** to share a chassis. The same loop — **scan** → **model** → **stress-test** → **route to decision** — turns in Area 2's debt fan charts and Area 25's flood maps alike. The same data backbone (Area 31) feeds them all; the same human futures-literacy (Area 32) decides whether anyone acts on what they say. The same fleet of agents widens the aperture in every domain. And the same three horizons — near, medium, long — recur, decision for decision, whether the subject is a winter adequacy margin or a fifty-year settlement pattern. The deep regularities are not a coincidence of style. They are the thesis.

Once you see the regularities, the prioritization stops looking like a ranked list of worthy projects and starts looking like what it is: a **sequencing instruction** for one build. You do not staff thirty-two teams. You staff one engine and point it, in turn, at the macro spine, at competitiveness, at society, and at the country's planetary and security boundaries. The five patterns on the facing pages are the load-bearing structure under the whole catalog.

## 1

shared loop — scan, model, stress-test, route — in every one of the 32 areas

## 2

load-bearing enablers under everything: the data backbone (31) and futures-literacy (32)

## 3

recurring horizons — near 0–2y, medium 2–6y, long 6–15y+ — in every domain

## 5

cross-cutting patterns that make 32 briefs one capability

### THE REFRAME

A country starting from zero does not buy thirty-two capabilities. It builds **one** — the engine, its data spine, and the people who can use it — and then extends it, tier by tier, at a marginal cost that falls with every domain added. The catalog is the order of operations.



*The same engine — scan, model, stress-test, route to decision, repeat — runs in every one; only the domain changes.*

ENSI · THE THREAD THROUGH ALL 32

ANGLE 04 · THE PATTERN MATRIX

# Five patterns, thirty-two domains: the same regularities recur everywhere.

If foresight were thirty-two unrelated projects, no single property would hold across them. It does. Map the five cross-cutting patterns against the four tiers and the cells fill in — not by accident, but because each area is the same engine wearing a different domain's clothes.

FIGURE B4.1 – PATTERN MATRIX · THE 32 AREAS × THE 5 RECURRING PATTERNS

	THE LOOP SCAN · MODEL · TEST · ROUTE	BACKBONE DATA (31) + LIT (32)	AGENT FLEET (32) ARCHETYPES	3 HORIZONS NEAR · MED · LONG	OPTION VALUE PAY
<b>Tier I · The Foresight Core</b> Areas 1–8 · centre, fiscal, growth, tech, labour, demography, energy, risk	●	●	●	●	●
<b>Tier II · Productive Capacity</b> Areas 9–16 · R&D, trade, digital, finance, infra, materials, firms, FDI	●	●	●	●	●
<b>Tier III · Society &amp; Quality of Life</b> Areas 17–24 · health, education, ageing, migration, cohesion, region, housing, wellbeing	●	●	●	●	●
<b>Tier IV · Environment, Security &amp; Enablers</b> Areas 25–32 · climate, food/water, nature, defence, cyber, regulation, data (31), literacy (32)	●	●	●	●	●

**Every cell is marked.** The five patterns are present in all four tiers because they are properties of the engine, not of any one domain. Areas 31 and 32 (bottom-right tier) are the backbone column made explicit — they appear as a *row* only because the catalog has to place them somewhere; functionally they sit *under* every other row.

**PATTERN 1 · THE LOOP**

Scan → model → stress-test → route.

Scanning agents widen the aperture; models turn signals into futures; red-teams and wind-tunnelling break the plan; routing puts the result into a budget, a law, a procurement. Every brief in the catalog is this loop, instantiated.

**PATTERN 2 · THE BACKBONE**

Data (Area 31) and futures-literacy (Area 32) are load-bearing. A scenario is only as good as the indicators beneath it; a brilliant scenario dies if no one in the room can use it. Remove either and all thirty other areas degrade at once.

ANGLE 04 · THE REMAINING REGULARITIES

# Agents dissolve the bottleneck; the horizons recur; the payoff is always option value.

## Pattern 3 — agents dissolve the attention bottleneck


Classical foresight is rationed by human attention. A handful of analysts can scan only so many signals, build only so many scenarios, and refresh them only so often — so foresight arrives as an occasional, expensive set-piece, a workshop every quarter. That scarcity is the reason most governments treat the future as a rumor: not because anticipation is hard, but because it was **expensive to do continuously**.

A standing fleet of agents removes the constraint. Scanning agents read the global information space every day; scenario agents spin and stress-test hundreds of futures; simulation agents run agent-based models of an economy or an epidemic; briefing agents deliver the relevant slice to the right desk at the moment of decision. Foresight stops being a quarterly artifact and becomes a continuous nervous system — the shift Angle 6 develops in full.

## Pattern 4 — the same three horizons, everywhere

Open any brief and the horizons line up: a **near** band (0–2 years) of operational moves already in play; a **medium** band (2–6 years) where curricula, reforms and permits are decided; and a **long** band (6–15 years and beyond) where capital, infrastructure and structural identity are committed. The classic policy error is confusing them — funding short retraining for a structural shift, or planning grand reform for a cyclical dip.

FIGURE B4.2 — ONE PAYOFF, REPEATED



**Pattern 5.** In every area the arithmetic is identical: a modest standing cost to see a shift early, set against the large, lumpy cost of being surprised by it late.

## Pattern 5 — the payoff is always option value

The countries that built this capability did not get better forecasts — forecasts are usually wrong. What they got was the ability to see a shift early, to have already rehearsed the response, and to **act while action is still cheap**. That is option value, and it is the payoff in every single brief: a skills gap caught before it becomes mass unemployment, a risk caught before it becomes a crisis, an asset caught before it becomes a stranded liability.

THE COMPOUND

A single area's option value is modest. The thesis is the **sum**: the same gap, captured across thirty-two domains and compounded across decades, is the difference between a country carried by the future and one that authors it. You do not buy that thirty-two times — you build the engine once and inherit the advantage everywhere it points.

## ANGLE 05 · THE MINIMUM VIABLE FORESIGHT STATE

## The order is the strategy: build the engine first, then point it tier by tier.

The prioritization is not just a ranking of worthy domains — it is a build sequence. A country starting from zero should stand up the core engine and its data and human foundations first, then extend into the high-leverage macro functions, and only then into society and the planetary boundaries. Skip the foundations and you do not get a faster foresight state. You get isolated reports nobody acts on.

There is a tempting mistake available to every government that discovers foresight: to commission the exciting, visible work first — the climate-adaptation scenario, the pandemic playbook, the AI-and-jobs study — because those are the topics that make headlines and win ministerial enthusiasm. It is the wrong order, and it fails predictably. A society-tier or environment-tier study built without the engine has nowhere to route its conclusions, no shared data backbone to keep its indicators live, and no futures-literate officials to act on it. It becomes a beautiful document that gathers dust — **foresight theatre**, not foresight capability.

The dependency logic runs the other way. Four areas form the **minimum viable foresight state**: the national foresight capability at the centre of government (Area 1), national risk and crisis anticipation (Area 8), the data architecture and decision-intelligence platform (Area 31), and futures literacy and capability-building (Area 32). The engine, its early-warning organ, its data spine, and its people. Stand those up and every subsequent domain is a low-cost extension of a system that already works — the scenario machinery, the agent fleet, the routing into the budget and the legislative cycle are all already running; a new area just points them at a new subject. Attempt the same domain without the foundations and you are rebuilding the whole apparatus, badly, thirty-two times.

### START HERE

**Areas 1, 8, 31, 32 are the minimum viable foresight state.** A lean foresight directorate of eight to twelve at the cabinet centre; an always-on cross-hazard early-warning organ; a data-mesh-and-knowledge-graph backbone for three priority domains; and a futures-literacy curriculum with facilitators seeded across ministries. Everything else compounds on that base.

# 4

foundation areas form the minimum viable foresight state (1, 8, 31, 32)

# 8-12

people in the lean centre-of-government directorate, year one

Area 1, first moves

# 3

priority domains for the first data-backbone build (Area 31)

Area 31, first moves

# “

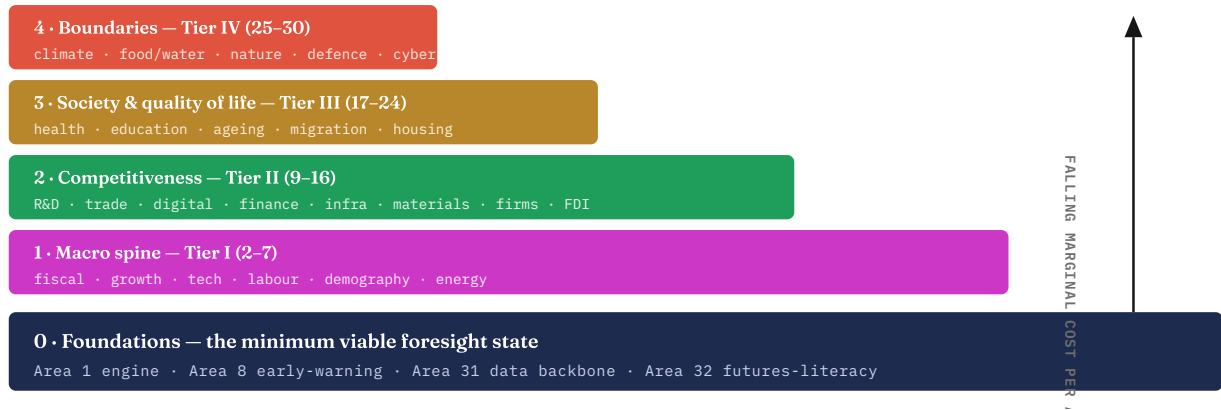
*Build the engine first and each new domain is a low-cost extension of a system that already works; attempt the society and environment tiers without the engine and you get isolated reports nobody acts on.*

ANGLE 05 · THE BUILD SEQUENCE

# Foundations, then the macro spine, then outward to the boundaries.

The four tiers are not parallel options to fund in any order. They are a dependency chain. Each stage inherits the engine, the data and the people from the stage before — and pays a falling marginal cost for every area it adds.

FIGURE B5.1 – THE BUILD SEQUENCE · DEPENDENCY FROM FOUNDATIONS OUTWARD



**Each stage stands on the one below.** The foundations carry the full width because every higher tier draws on them. A study at any upper stage that bypasses the base has nowhere to route, no live data, and no one trained to act — the signature of an orphaned report.

- 1 **Foundations are non-negotiable and come first.** Without the engine (1), the early-warning organ (8), the data backbone (31) and futures-literacy (32), nothing downstream routes, refreshes, or gets acted on.
- 2 **The macro spine has the highest leverage per unit of effort.** Fiscal, growth, technology, labour, demography and energy are where anticipation buys the most GDP and avoids the most crisis — extend here first.
- 3 **Competitiveness compounds on the spine.** R&D, trade and digital bets only make sense once the macro picture and the data backbone they depend on are live.
- 4 **Society and boundaries are extensions, not foundations.** They are where foresight most visibly touches lives — but built without the base they are the reports that gather dust.

**WHY THE ORDER PAYS**

The marginal cost of a new area **falls** as the sequence proceeds. By the time a state reaches the society and environment tiers, the scenario machinery, the agent fleet and the routing into budget and law are already running. A new domain is a configuration, not a construction project.

ANGLE 06 · THE AGENTIC SHIFT & THE GUARDRAILS

# Agents turn foresight from a quarterly artifact into a continuous nervous system.

This is the move that makes the document an ENSI document rather than a restatement of classical futures studies. Traditional foresight is bottlenecked by human attention, so it arrives as an occasional, expensive set-piece. A standing fleet of agents dissolves that bottleneck — and in doing so changes foresight from something a country *has* into something a country continuously *does*. But the corollary is sharp: when the machine never sleeps, discipline matters more, not less.

A handful of analysts can scan only so many signals, build only so many scenarios, and refresh them only so often. That is why, for most of its history, foresight has been a quarterly workshop and a binder on a shelf. The agentic layer breaks the constraint at its root. Agents can read the entire global information space every day, spin hundreds of scenarios, run agent-based simulations overnight, attack every plan, and deliver the relevant two pages to the right minister at the moment of decision. The set-piece becomes a **standing organ of the state** — always on, always current, embedded in the decisions that are being made anyway.

What follows are the five agent archetypes that recur, in different combinations, in all thirty-two areas. Each does something that scarce human analysts structurally never could — not because the humans were less capable, but because attention was always rationed and now it is not. The humans do not disappear; they move up the stack, from doing the scanning to judging what the scanning found.

## 5

agent archetypes recur across all 32 areas

**quarterly** → **continuous**

the shift from set-piece to standing nervous system

**humans up-stack**

from doing the scan to owning the judgement

**more, not less**

discipline required when the machine never sleeps

## “

*Foresight stops being a quarterly artifact and becomes a continuous, living nervous system for the state. Agents widen the aperture; humans keep the accountability.*

ENSI · THE AGENTIC ENGINE

## ANGLE 06 · THE FIVE ARCHETYPES

## What disciplined agents make possible that scarce humans never could.

The same five roles staff every area's engine. Read them as the standing crew of the foresight state — the watch, the imaginers, the modellers, the adversary, and the messenger.

### 01 Scanning agents

Ingest the entire global information space — preprints, patents, funding rounds, vacancy text, trade and sensor feeds — every day, and flag trend-breaks the moment they cross a threshold. **What humans never could:** read everything, continuously, and never tire of the boring 99% where the weak signal hides.

### 02 Scenario-generation agents

Recombine signals into fresh narrative branches, spin and stress-test hundreds of futures, and pressure-test the canonical scenarios for staleness. **What humans never could:** explore the combinatorial space of futures at a breadth no workshop can reach, then prune to the plausible.

### 03 Simulation & ABM agents

Run agent-based models of households, firms, regions, labour markets or an epidemic — testing how a policy or a shock propagates through the real system. **What humans never could:** run thousands of behavioural paths nightly and refresh the fan chart as the world moves.

### 04 Red-team agents

Adversarially attack every major plan, hunting the assumption that, if false, breaks it — and generate the compound shocks no one war-gamed. **What humans never could:** attack every plan, relentlessly, without the institutional reluctance to embarrass a colleague.

### 05 Monitoring, early-warning & briefing agents

Track the leading indicators that say which scenario the country is sliding into, fire calibrated alerts with lead time to act, and compress the whole picture into the two-page anticipatory brief the right decision-maker needs at the moment of choice — each tailored to the decisions that owner actually controls. **What humans never could:** hold the cross-hazard picture in one head, around the clock, and deliver its relevant slice to thirty desks at once without it going stale on the way.

#### THE DIVISION OF LABOUR

Agents widen the aperture; humans keep the accountability. The archetypes do the scanning, the spinning, the modelling and the attacking. The human foresight directorate **curates, judges and owns** what goes up — and remains answerable for every consequential call.

## ANGLE 06 · THE DISCIPLINE

# When the machine never sleeps, the guardrails are the capability.

An always-on agent fleet can be a continuous nervous system — or a continuous source of confident, unaccountable error at scale. The difference is not the model. It is the discipline wrapped around it. A foresight state that automates the scanning but not the governance has not built a capability; it has built a liability that briefs the prime minister.

The agentic shift does not relax the need for rigour — it raises it. A human analyst who makes a judgement makes one, slowly, with a name attached. A fleet of agents makes thousands, fast, and the temptation is to trust the output because it is fluent and arrived on time. That is exactly the failure mode to design against. The five guardrails below are the standing discipline that keeps an always-on system trustworthy — and they are non-optional precisely because the system is continuous.

- 1 **Human-in-the-loop on consequential calls.** Agents widen the aperture and draft the analysis; a named human curates, judges and owns anything that routes into a budget, a law or a crisis response. The accountability never leaves a person.
- 2 **Audit logs.** Every signal flagged, scenario spun and alert fired is logged and replayable — so a decision can be reconstructed, contested and learned from, not taken on faith because it looked plausible on the day.
- 3 **Provenance.** Every claim is traceable to source. A briefing that cannot show where its evidence came from is a rumor with better typography — and the whole point of foresight is to refuse to govern on rumor.
- 4 **Evaluation.** The agents are themselves measured — calibration of their alerts, hit-and-miss rates on materialised risks, drift in their scenarios — so the fleet is held to a standard, not assumed to be right because it is automated.
- 5 **An external test the agents cannot fake.** The system is checked against reality it does not author — materialised events, independent data, a red-team outside the fleet — so it cannot quietly grade its own homework.

## THE GUARDRAIL MODEL — NON-NEGOTIABLE

Continuous foresight without provenance, audit and an external check is not a faster capability — it is automated overconfidence with a direct line to power. **Discipline is what separates a nervous system from a hallucination.** The agents earn their place only inside the guardrails.

“

*Agents change foresight from a quarterly set-piece into a continuous nervous system — and that is exactly why the discipline has to be tighter, not looser. The imagination and the judgement stay human.*

ENSI · THE AGENTIC SHIFT, AND ITS PRICE

ANGLE 07 · SEGMENTATION – WHERE TO START

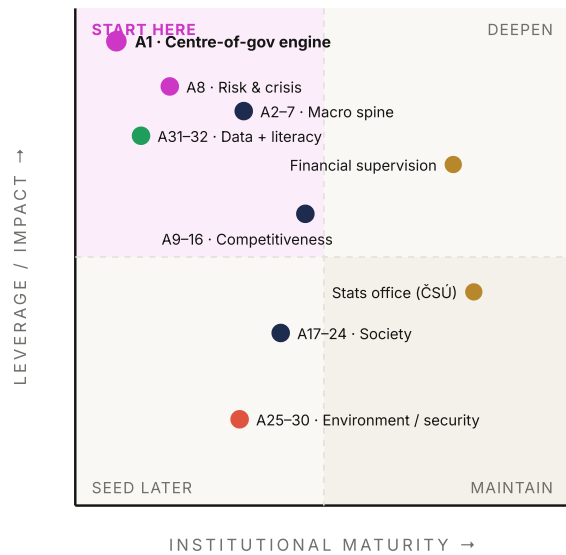
# Not every domain deserves your first foresight dollar

The 32 areas are not equal candidates for early investment. The first move belongs where two conditions meet: **leverage is highest** — getting foresight right there changes everything downstream — and **institutional maturity is lowest**, so the standing capability does not yet exist to be improved. Plotted on those two axes, the four tiers stop being a ranking and become a map of where the marginal foresight euro buys the most future.

The catalog ranks the 32 areas by *leverage* alone, and that is the right ordering for a reference layer. But leverage is only half of a *start-here* decision. A domain can be enormously consequential and already well-served — a country with a strong central bank does not need to stand up financial-stability foresight from scratch. The decision rule is the gap between how much the domain matters and how much capability is already in place to anticipate within it.

The 2x2 below makes that gap visible. The vertical axis is **leverage / impact**: how far a foresight advantage in this domain propagates into every decision taken downstream. The horizontal axis is **institutional maturity / readiness**: how much standing anticipatory capacity the typical mid-sized European state already has in the domain. The first foresight dollar belongs in the **top-left quadrant** — high leverage, low maturity — because that is where the same money moves the most and meets the least existing structure.

FIG. 7.1 – LEVERAGE × INSTITUTIONAL MATURITY



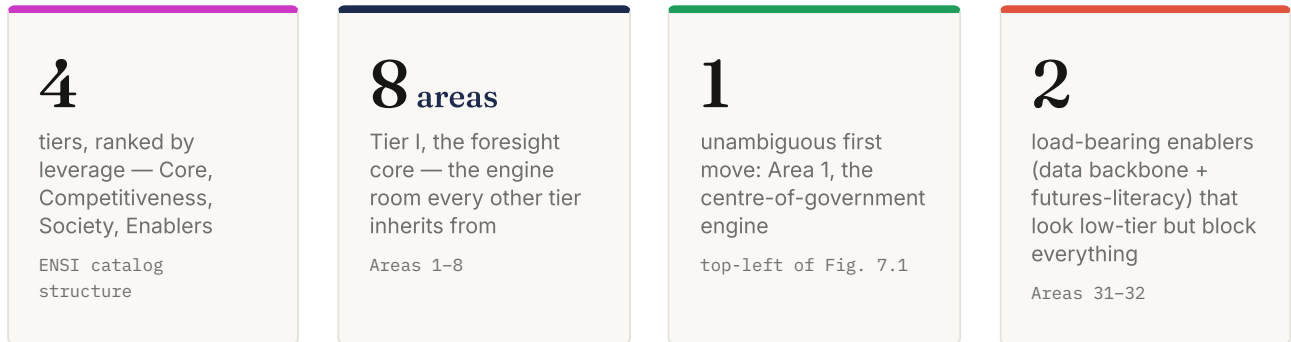
**Read the top-left quadrant first.** Positions are indicative of relative leverage and the standing capacity a typical mid-sized EU state already holds, not precise scores. Magenta = the unambiguous first move.

**THE START-HERE RULE**

Build where the **leverage curve is steepest and the maturity curve is flattest**. Area 1 — the national foresight capability at the centre of government — sits at the extreme top-left: maximum leverage (every other area runs on it), near-zero maturity (almost no state has a standing central engine). It is the unambiguous first move; everything else plugs into it.

ANGLE 07 · SEGMENTATION – CONTINUED

## Where each tier ranks — and the first move it earns



TIER	WHY IT RANKS HERE	THE FIRST MOVE
<p><b>Tier I — Core (1–8)</b> The engine room</p>	<p>Highest leverage of all: the central capability plus the macro-fiscal, growth, technology, labour, demographic, energy and risk spine. Get these right and every other tier inherits the advantage; get them wrong and no downstream foresight can compensate.</p>	<p>Stand up <b>Area 1</b> at the cabinet centre — a Chief Foresight Officer, a lean directorate, the national scenario set and signals registry. Then wire Areas 2 and 8 (fiscal, risk) into the budget and crisis machinery.</p>
<p><b>Tier II — Competitiveness (9–16)</b> Climbing the ladder</p>	<p>Determines whether the economy keeps climbing the complexity ladder — innovation, trade, digital, finance, infrastructure, materials, the firm base, investment. High leverage, but most states already hold fragments (a statistics office, a central bank, an investment agency).</p>	<p>Connect, don't rebuild: anchor industrial-strategy and supply-chain foresight (Areas 3, 10) to the centre's shared scenario set so national bets compound rather than contradict.</p>
<p><b>Tier III — Society (17–24)</b> How people live</p>	<p>Health, education, ageing, migration, cohesion, regional balance, housing, wellbeing. Leverage is real but the time-constants are long and the existing institutions (schools, hospitals, pension funds) are dense — foresight here deepens an established system rather than creating one.</p>	<p>Seed selectively where a demographic certainty is already locked in — pension and long-term-care sustainability, the depopulating-region early-warning map — using the centre's shared population baseline.</p>
<p><b>Tier IV — Enablers (25–32)</b> Boundaries + backbone</p>	<p>Climate, food/water, nature, defence, cyber, regulation — and, load-bearing for everything, the <b>data architecture (31)</b> and <b>futures-literacy (32)</b>. These two look like the bottom of the list but are the substrate the whole system runs on: a scenario is only as credible as the data feeding it.</p>	<p>Treat data architecture as strategic infrastructure from day one — stand up the data-mesh for three priority domains in parallel with Area 1, not after. Defence/cyber foresight wires into the national-security secretariat.</p>

THE SEQUENCING INSIGHT

The catalog ranks by leverage; the **build order** follows leverage *minus* maturity. That is why the minimum viable foresight state is not "Tier I in numerical order" but a diagonal cut across the tiers — Area 1 (the engine), Area 8 (the early-warning system), and Areas 31–32 (the data backbone and the literacy to use it). The order is the strategy: build the engine first, lay the data rails beside it, and every one of the remaining areas plugs into a capability that already exists.

## ANGLE 08 · BENCHMARKING – THE EXEMPLARS

# The leading foresight states didn't get better forecasts

They built **standing capability** and wired it into decisions. Singapore, Finland, the UAE, Canada, the EU and the UK do not out-predict everyone else — forecasts are usually wrong everywhere. What separates them is structural: a permanent unit, a fixed cadence, and a hard route from the foresight output into the budget, the legislative programme or the national strategy. The lesson is never the method. It is the **plumbing**.

Each exemplar below is profiled on three things, because those three are what actually transfers across borders: *what they built* (the standing institution), *how it routes into decisions* (the wiring that stops a report dying on a shelf), and *the one lesson* a country starting today should copy. The institutions and their signature outputs are drawn from ENSI's source library of the world's leading foresight bodies.

## Singapore — Centre for Strategic Futures

**Built.** The CSF, inside the Strategy Group of the Prime Minister's Office, and the **Scenario Planning Plus / RAHS** toolkit — a standing futures function at the very centre of the state, now past its fifteenth year.

**Routes in.** Foresight feeds directly into whole-of-government strategic planning; the biennial *Foresight* journal and scenario sets are shared infrastructure that ministries plan against, not advisory papers.

**The one lesson.** Put it at the centre, keep it permanent, and make the scenario set a shared asset the whole of government uses — once, well, rather than thirty times badly.

## Finland — Sitra + the Committee for the Future

**Built.** A three-part system: **Sitra** (the independent futures fund and its Megatrends), the parliamentary **Committee for the Future**, and the statutory **Government Report on the Future** each electoral term.

**Routes in.** The Government must table a foresight report and Parliament's own committee must respond — foresight is written into the constitutional rhythm of governing, not bolted on.

**The one lesson.** A **statutory hook** beats good intentions. When law requires a future-facing report and a parliamentary answer, foresight survives the minister who championed it.

## UAE — Dubai Future Foundation

**Built.** A sovereign futures agency (DFF) with its own mandate, the Museum of the Future as a public anchor, and signature outputs such as *The Global 50* future-opportunities report.

**Routes in.** Foresight is fused with delivery — opportunity-scanning flows into national strategies, accelerators and procurement, treating the future as a pipeline of bets to place, not a forecast to admire.

**The one lesson.** Foresight gains political force when it is **visible and opportunity-framed** — anchored in institutions citizens and ministers can see, not buried in a secretariat.

## Canada — Policy Horizons

**Built.** **Policy Horizons Canada**, the federal government's dedicated foresight organization, with a repeatable annual cadence — its flagship *Disruptions on the Horizon* scan assesses dozens of potential disruptions across society, economy, environment and geopolitics.

**Routes in.** A standing methodology and an annual product give departments a common, refreshed disruption map to plan against, with training that spreads futures skills across the public service.

**The one lesson.** **Cadence and method** matter more than any single brilliant report — a repeatable annual scan that everyone reads beats an occasional masterpiece nobody acts on.

ANGLE 08 · BENCHMARKING — THE SUPRANATIONAL TIER

## Two more exemplars — and the pattern they share

### EU — JRC / ESPAS & the Strategic Foresight Report

**Built.** An institutionalized cycle: the **Joint Research Centre's** foresight capacity, the inter-institutional **ESPAS** process, and an annual **Strategic Foresight Report** from the Commission (e.g. the 2023 edition on sustainability and open strategic autonomy).

**Routes in.** The annual report feeds the Commission's work programme and concrete actions — foresight is embedded in the Union's policy-making cycle, not a side study.

**The one lesson.** An **annual, named report tied to the work programme** turns foresight into a recurring obligation of governing — and gives member states a ready ecosystem to nest inside.

### UK — Government Office for Science

**Built.** **GO-Science** and its flagship Foresight programme, which run structured evidence reviews and "future worlds" exercises (e.g. the Future of Mobility, looking out to 2040), alongside the national risk-assessment practice.

**Routes in.** Foresight projects are commissioned to inform specific cross-departmental decisions, with a chief-scientific-adviser network carrying the evidence into each ministry.

**The one lesson.** **Tie every foresight project to a live decision** and give it a named official who owns the route in — foresight that is commissioned by a decision rarely dies orphaned.

None of these states bought a better crystal ball. They bought a standing function, a fixed cadence, and a statutory or institutional route from the scenario set into the decision — and that route, not the forecast, is the asset.

### ENSI Foresight Division · the benchmarking finding

STATE	THE STANDING INSTITUTION	HOW IT ROUTES INTO DECISIONS	THE ONE LESSON
Singapore	Centre for Strategic Futures (PMO); Scenario Planning Plus / RAHS	Whole-of-government scenario set; biennial <i>Foresight</i> journal as shared infrastructure	Put it at the centre, keep it permanent, share the scenarios
Finland	Sitra + parliamentary Committee for the Future	Statutory Government Report on the Future + a parliamentary answer each term	A statutory hook outlasts the champion
UAE	Dubai Future Foundation (sovereign agency)	Opportunity-scanning fused with strategy, accelerators and procurement	Make foresight visible and opportunity-framed
Canada	Policy Horizons Canada	Annual <i>Disruptions on the Horizon</i> scan + futures training across the service	Cadence and method beat the occasional masterpiece
EU	JRC / ESPAS + Strategic Foresight Report	Annual report feeds the Commission work programme & concrete actions	A named annual report becomes an obligation of governing
UK	Government Office for Science — Foresight programme	Decision-commissioned reviews carried in by the CSA network	Tie every project to a live decision and a named owner

#### THE COMMON DENOMINATOR

Across six very different political systems, the variable that predicts impact is not analytical sophistication. It is whether there is a **standing home**, a **fixed cadence**, and a **hard wire** into a decision that is being made anyway. ENSI's agentic engine is the modern version of the same idea: it industrializes the scanning and scenario work so the standing function can run continuously rather than as a set-piece — but the wiring into the decision still has to be built by hand.

ANGLE 08 · BENCHMARKING – THE HONEST MIRROR

## Where a country like Czechia stands

Czechia has real foresight ingredients — they are simply not yet assembled into an engine. The honest reading is neither "we have nothing" nor "we are fine": it is that the country holds capable parts and a rich external ecosystem, but **no standing central foresight function** to run them once for the whole of government.

**What is already there.** A capable national statistical office — the **Czech Statistical Office (ČSÚ)** — already engaged with the UNECE HLG-MOS modernisation agenda, giving the country a credible data backbone to build on. Sectoral pockets of anticipation exist: sector-skills councils on the labour side, the R&D council (RVVI) and Technology Agency on the science side, a digital agency, and a transmission-system operator that already models energy adequacy. And Czechia sits inside a dense **EU foresight ecosystem** — the JRC, ESPAS, the annual Strategic Foresight Report, Horizon Europe — that supplies scenarios, methods and signals it can nest inside rather than rebuild.

**What is missing.** There is no standing central foresight engine — no unit at the cabinet centre that owns a shared national scenario set, runs the scanning machinery once for the whole government, and holds line ministries to wind-tunnelling their plans against more than one future. The capable parts therefore operate as thirty private spreadsheets, not one shared spine. As a frontline-adjacent NATO and EU member with limited fiscal slack and a sharp demographic transition already locked in, the asymmetry is unforgiving: Czechia cannot afford to buy its way out of a crisis it failed to imagine.

### The Czech foresight balance sheet

- ▶ **Asset** — ČSÚ, a capable statistical office on the UNECE modernisation track
- ▶ **Asset** — embedded in the EU foresight ecosystem (JRC, ESPAS, Strategic Foresight Report, Horizon Europe)
- ▶ **Asset** — sectoral anticipation pockets: sector-skills councils, RVVI / Technology Agency, the digital agency, the TSO's adequacy models
- ▶ **Gap** — no standing central foresight unit at the cabinet centre
- ▶ **Gap** — no shared national scenario set the whole government plans against
- ▶ **Gap** — no statutory hook making foresight a required input to budget and legislation

#### THE OPPORTUNITY

The exemplars show the gap is bridgeable cheaply. Finland's statutory report, Canada's annual cadence and Singapore's shared scenario set are all **organisational moves, not megaprojects**. Czechia's parts are good enough that the decisive act is assembly — stand up Area 1 at the centre and connect what already exists.

1

capable national statistical office (ČSÚ) as a ready data backbone

0

standing central foresight engines today — the decisive gap

4+

sectoral anticipation pockets to connect, not rebuild

EU

foresight ecosystem to nest inside (JRC · ESPAS · SFR · Horizon Europe)

ANGLE 09 · FAILURE MODES – THE ANTI-PATTERNS

# Most foresight fails at the plumbing, not the prediction

The report is written and nobody acts. That single sentence accounts for more dead foresight units than any analytical error ever did. Foresight rarely fails because the scan missed the signal or the scenario was wrong; it fails because there was **no route from the output to a decision** — no owner, no cadence, no mandate, no data to refresh on. The failures are organisational, and so are the fixes.

Six failure modes recur across the institutions that tried foresight and lost it. Each has a predictable cause and a fix that ties straight back to the three things this report argues a foresight capability must be: an **engine** rather than a report, kept current by an **agentic refresh**, and owned at the **centre of government (Area 1)** with a hard route into the decisions being made anyway.

FAILURE MODE	WHY IT HAPPENS	THE FIX – AND WHERE IT TIES BACK
<b>Orphaned reports</b> written, shelved, forgotten	The output has no pre-agreed destination — no decision, budget line or law it feeds. A brilliant scenario set lands on a desk with no obligation attached, and the next crisis quietly contradicts it.	Wire every product to a <b>live decision before it is written</b> (the UK GO-Science discipline). The centre-of-government engine (Area 1) owns a statutory hook into budget and legislative cycles, so foresight is a required input, not a courtesy.
<b>No data backbone</b> scenarios built on sand	A scenario is only as credible as the indicators feeding it, and an early-warning system blind across domains cannot see where modern crises form. Without integrated data, the unit improvises from private spreadsheets.	Treat data architecture ( <b>Area 31</b> ) as strategic infrastructure from day one — the data-mesh and signals registry built <i>beside</i> the engine, so every scenario refreshes from one shared, reusable spine.
<b>No mandate / wrong altitude</b> real but powerless	The unit is buried in a line ministry with no convening power, or it has authority but sits too far from the budget. It cannot make other departments wind-tunnel their plans, so its work is advisory at best.	Place it at the <b>cabinet centre</b> reporting to the cabinet secretary, with embedded ministry liaisons (Area 1). Finland's statutory route shows mandate is what outlasts the champion who created the unit.
<b>Captured by short-termism</b> the electoral clock wins	Ministers are rewarded for what shows up inside a four-year cycle; the long, certain forces — ageing, the energy transition, machine intelligence — are treated as someone else's problem until they arrive as a crisis.	Hard-wire the <b>long horizons into the calendar</b> : a statutory future report each term and a standing weekly anticipatory brief make the 25-year view a recurring obligation, not a discretionary luxury.
<b>One-off exercise</b> no continuous refresh	Foresight arrives as an occasional, expensive set-piece — a workshop, a glossy report — then goes stale because no one can afford to run it again. The world moves; the scenarios do not.	Replace the set-piece with the <b>agentic refresh</b> : a standing fleet of scanning, scenario and red-team agents keeps the scenario set live and fires when a signal crosses a threshold. Cadence, as Canada shows, beats the occasional masterpiece.
<b>Foresight as theatre</b> signalling, not deciding	The unit exists to be seen — a museum, a launch event, a strategy document — with no obligation that anything change as a result. Activity substitutes for influence, and the theatre is defunded the moment budgets tighten.	Bind every exercise to an <b>owned decision and a metric</b> : share of major decisions formally wind-tunnelled, lead time from signal to response. The agentic engine industrialises the work so the unit earns its keep in decisions, not appearances.

ANGLE 09 · FAILURE MODES – THE THROUGH-LINE

### Every fix routes back to the same three things

Read the failure table as a column, not six rows, and one pattern resolves: every fix ties back to the same three structural commitments this report has argued from the first page. **The engine** — foresight as a standing capability, not a report — kills the orphaned-report and one-off-exercise failures. **The agentic refresh** — a continuous fleet of scanning, scenario and red-team agents — kills staleness and theatre by making the work cheap enough to run forever. And **Area 1 ownership at the centre of government** — with a statutory hook into budget and legislation — kills the mandate, altitude and short-termism failures by giving foresight a home with the authority to be acted on.

This is why the diagnosis matters more than it first appears. A country that believes its foresight failed because the forecast was wrong will respond by buying better analysis — and watch the next unit die the same death. A country that understands the failure is *plumbing* will instead invest in the route from output to decision: the owner, the cadence, the mandate, the data backbone. The cheaper, more durable fix is almost always the organisational one.

**THE WARNING**

The single most expensive mistake in foresight is to fund the analysis and starve the wiring. A scanning fleet with no route into the budget is theatre with a dashboard. **Build the route first, then turn on the engine** — never the reverse.

#### The three through-lines

- 1 The engine, not the report.** A standing function that scans, models, stress-tests and routes — continuously. Fixes orphaned reports and one-off exercises.
- 2 The agentic refresh.** A fleet of scanning, scenario, simulation and red-team agents keeps the picture live and fires on threshold-crossing signals. Fixes staleness and theatre.
- 3 Area 1 ownership.** A home at the cabinet centre with a statutory hook into budget and legislation. Fixes wrong altitude, missing mandate and capture by short-termism.

---

<p><b>6</b></p> <p>recurring failure modes — all organisational, none analytical</p>	<p><b>3</b></p> <p>structural fixes they all reduce to: engine · agentic refresh · Area 1</p>
<p><b>1</b></p> <p>sentence that kills most units: "the report is written and nobody acts"</p>	

---

Foresight does not die at the moment of prediction. It dies in the gap between a finished report and a decision no one obliged anyone to make — and that gap is closed with plumbing, not with a better forecast.

ENSI Foresight Division · **the failure-modes finding**

TIER I · AREAS 1-8

# The Foresight Core

The engine room and the macro spine. These eight functions are the central capability itself plus the fiscal, growth, technology, labour, demographic, energy and risk watch every other decision hangs from. Get them right and every later tier inherits the advantage; get them wrong and no downstream foresight can compensate.

1 Foresight Capability at the Centre

2 Macro-Fiscal & Public Finance

3 Growth & Industrial Strategy

4 Technology & Emerging Tech

5 Labour Market & Skills

6 Demographic & Population

7 Energy System & Security

8 National Risk & Crisis

# 01 The Foresight Capability at the Centre of Government

A permanent strategic-foresight unit in the cabinet centre that routes anticipatory intelligence straight into the budget, the legislative programme and the security review — infrastructure, not advice.

**Why it ranks first.** Every other area runs on this one: absent a home, a method and a route into power, foresight is reinvented ad hoc, politicised, or defunded the moment its champion moves on. The payoff is not better forecasts — governments are mediocre forecasters — but institutional reflexes: the habit of asking *"and if the world turns out otherwise?"* before money is committed.

## 5 · 15 · 30

HORIZON · stress-test programme & budget / shape strategy & capital bets / interrogate the structural givens

## Signals registry

SIGNAL · one curated, deduplicated weak-signal feed for the whole of government — not thirty private spreadsheets

## Always-on fabric

AGENT · scanning, scenario-generation, simulation, red-team & briefing agents widen the aperture; humans keep the accountability

### FIRST MOVE · 0-12 MONTHS

Recruit a lean directorate of 8-12 under a Chief Foresight Officer reporting to the cabinet secretary; publish the first national scenario set, stand up the signals registry, embed liaisons in finance/economy/interior, and deliver the first wind-tunnelled budget input. **Build the engine first — the other thirty areas plug into it.**

# 02 Macro-Fiscal & Public-Finance Foresight

Budgeting for the futures that have not happened yet — scenario-based budgets, debt stress-tested against shocks, and a live register of the contingent liabilities priced before they crystallise.

**Why it ranks here.** Fiscal space is the master constraint: it decides whether a country can invest, absorb a shock, or fund the transitions every other area demands. The binding risks are almost never in the headline deficit — they are the off-balance-sheet pension promises, guarantees and care bills that reactive budgeting discovers only at the moment of crisis, when options are worst and borrowing costs highest.

## 5 · 15 · 30

HORIZON · medium-term framework robustness / structural-reform timing / debt under full demographic + climate transition

## Second derivative

SIGNAL · watch not the deficit but the rate at which structural pressures accelerate — yields, demographics, contingent-liability triggers

## Nightly Monte Carlo

AGENT · simulation agents refresh the debt fan chart as the world moves; red-team agents hunt the shock combo no one war-gamed



*The demographic spending wave is the single most predictable — and most under-provisioned — fiscal event of the next three decades.*

ENSI FORESIGHT DIVISION · AREA 2, MACRO-FISCAL FORESIGHT

## 03 Economic Growth & Industrial-Strategy Foresight

Treating growth as navigation on the map of what a country can plausibly become next — and using anticipation to place scarce industrial and innovation bets on the right rungs of the complexity ladder.

**Why it ranks here.** Long-run prosperity is a story of structural transformation, not of doing the same things more efficiently. Growth is path-dependent: a country can only easily move into products "nearby" in capability space to what it already makes — so diversification is a foresight problem, defined by where global demand and technology will be in fifteen years, not today. Pick the wrong rung and you waste a decade of subsidy.

### 5 · 15 · 30

HORIZON · which programmes to fund / which capability platforms to build / the structural identity of the economy

### Relatedness density

SIGNAL · is the country accumulating the adjacent capabilities that make a high-complexity bet feasible? (Atlas of Economic Complexity, product space)

### Live product space

AGENT · complexity-modelling agents recompute reachable high-value rungs as data lands; red-team agents test the adjacency

#### FIRST MOVE · 0-12 MONTHS

Build the national economic-complexity dashboard, run a relatedness analysis to shortlist reachable high-value targets, scenario-test the top anchor value chains, and deliver a first prioritised diversification map to the economy minister.

**Separate feasible strategy from nostalgia dressed as strategy.**

## 04 Technology & Emerging-Tech Foresight

A continuous, structured watch on emerging technologies — and on their second- and third-order consequences — to detect disruption while it is still cheap and reversible to act on.

**Why it ranks here.** Technology is the fastest-moving force on growth, security and quality of life, and it respects no electoral calendar. The cost of acting rises steeply along the diffusion curve: intervene early on standards, skills and regulation and you shape the trajectory for almost nothing; intervene late and you buy influence over an entrenched system at ruinous cost. The expensive surprises are the second-order effects — what AI does to labour, what cheap biotech does to biosecurity, what quantum does to encryption.

### 5 · 15 · 30

HORIZON · skills, sandboxes & procurement / research priorities & standards / transformative AGI, synthetic biology, fusion

### Diffusion velocity

SIGNAL · track not just maturity but how fast a capability spreads and which downstream systems it perturbs (patents, preprints, compute benchmarks)

### Futures-Wheel agents

AGENT · auto-generate 2nd- and 3rd-order consequence trees for each signal; early-warning agents fire at watch thresholds

#### FIRST MOVE · 0-12 MONTHS

Stand up the emerging-tech register and scanning fabric; run Futures-Wheel assessments on the three most consequential technologies — frontier AI first; establish a regulatory-sandbox and skills fast-track; deliver a first quarterly emerging-tech brief to the cabinet.

## 05 Labour Market & Skills Foresight

Skills are the slowest variable in any economy. Foresight treats the workforce as something a country builds a decade ahead of demand — turning automation panic into managed transition and demographic change into a wage dividend.

**Why it ranks here.** The real threat is skills *mismatch*, not job destruction — and mismatch is forecastable. The automation debate fixates on net job counts, where it is wrong in both directions; the violent churn happens underneath stable totals, and the displaced are rarely the people the new roles need. Generative AI is the first wave to expose cognitive and white-collar work at once, compressing transitions that once took a generation into years.

### 3 · 7 · 15

HORIZON · where ALMP spend moves first / what universities & VET expand or close / which new occupational families to seed now

### Vacancy text

SIGNAL · online job postings shift weeks ahead of hiring stats; the quiet appearance of new tool requirements inside old titles (Cedefop OVATE, Lightcast)

### Exposure-mapping

AGENT · agents keep the occupation-by-task risk map current as new AI capabilities ship; ABM simulates worker transitions

#### FIRST MOVE · 0-12 MONTHS

Stand up the live vacancy-and-exposure dashboard; publish a first national skills outlook to three horizons; run one backcasting exercise on a strategic workforce — energy, care or AI-orchestration — with a costed pipeline the ministries commit to fund.

## 06 Demographic & Population Foresight

Demography is the most predictable force shaping a nation and the most chronically ignored. The job is forcing slow, compounding certainties into decisions taken on fast political clocks.

**Why it ranks here.** Demography's danger is its very predictability: because nothing happens suddenly, nothing feels urgent, and the cheap early window closes unnoticed. There is no fast fix — you cannot conjure forty-year-olds, and pro-natal or migration responses take a generation to register. Most of Europe is at the inflection where post-war cohorts retire en masse as small younger cohorts enter work; the choices that soften the next thirty years must be made this decade.

### 10 · 25 · 50

HORIZON · school & care capacity / pension & support-ratio solvency / settlement & housing stock locked in for half a century

### Births & cohorts

SIGNAL · births forecast school and labour entry with near-perfect lead time; the desired-vs-actual fertility gap, emigration intentions of the young and skilled

### One baseline

AGENT · cross-domain agents propagate one coherent population baseline into the skills, health & energy areas — no contradictory assumptions

#### FIRST MOVE · 0-12 MONTHS

Publish a single authoritative population baseline every department must plan against; run fiscal-sustainability stress tests on pensions and long-term care under three scenarios; produce a regional early-warning map of areas approaching service-viability thresholds, with a costed response menu.

## 07 Energy System & Security Foresight

Energy sits at the intersection of growth, security and climate, and the next two decades rebuild it end to end. This is where getting foresight wrong is most expensive — and most visible.

**Why it ranks here.** Security and decarbonisation are not opposites to be balanced but a single system to be *sequenced* — and sequencing is a foresight problem. Energy assets have 20-to-40-year lifespans, so a transmission line or nuclear unit commissioned today locks in cost or risk for decades. Electrification of transport and heat is about to surge demand precisely as variable renewables rise and thermal plant retires; the reliability margin thins before it improves, and that vulnerable middle passage must be planned through, not stumbled into.

### 3 · 10 · 30

HORIZON · winter adequacy & reserve / build-out & coal/nuclear retirement / deep bets — new nuclear, hydrogen, long-duration storage

### Adequacy + security

SIGNAL · fuse grid, market and geopolitical feeds into one picture — reserve margins, storage, prices, mineral chokepoints (ENTSO-E, IEA)

### Compound-shock

AGENT · red-team agents build cold-snap + outage + supply-cut scenarios and probe where the system breaks before winter does

#### FIRST MOVE · 0-12 MONTHS

Build the integrated adequacy-and-security dashboard; run a full scenario set to 2050 with explicit security stress tests; wind-tunnel the current generation and grid investment plan — surfacing the no-regret moves and the bets that need real-options staging.

## 08 National Risk & Crisis Anticipation

Almost every state keeps a national risk register and treats it as a document — reviewed annually, shelved, contradicted by the next surprise. This reimagines it as a living, always-on early-warning system.

**Why it ranks here.** Real crises are interconnected and arrive faster than annual review cycles can track. Anticipation beats reaction more starkly here than anywhere: in a crisis, a day of early warning is worth enormous amounts and decays fast, while the cost of preparedness is modest and the cost of surprise catastrophic. The capability gap is between the speed of modern shocks and the cadence of the institutions meant to see them coming.

### Days → 20yr

HORIZON · uniquely spans every horizon at once — what is escalating now / which risks are intensifying / which structural threats to build resilience against

### The correlations

SIGNAL · weak signals live at the edges — a tremor in one system that historically precedes failure in another; the cascade siloed watch-floors miss by construction

### Correlation agents

AGENT · scan health, financial, cyber, supply, climate & geopolitical streams round the clock and hunt cross-domain cascade precursors

#### FIRST MOVE · 0-12 MONTHS

Stand up the always-on cross-hazard scanning and fusion capability; rebuild the national risk assessment as a living, cascade-based system rather than a static register; run a compound-shock exercise stress-testing two or three simultaneous crises — then fix the gaps it exposes.

TIER II · AREAS 9-16

# Productive Capacity & Competitiveness

Whether the economy keeps climbing the complexity ladder: where to point the science budget, how to read trade and capital flows before chokepoints close, what digital and physical backbone to provision ahead of demand, and how to see financial trouble forming.

9 Innovation & R&D

10 Trade & Geoeconomics

11 Digital Infrastructure

12 Financial & Systemic Risk

13 Infrastructure & Capital

14 Critical Minerals

15 SME & Firm Dynamics

16 Foreign Investment

## 09 Innovation & R&D Priority-Setting

Pointing the science budget at tomorrow's strengths — funding the capabilities that compound into national edge fifteen years out, not the fields that produced last decade's citations.

### WHY IT RANKS

R&D is the slowest-acting and highest-multiplier lever a state holds: a misallocated decade of research wastes a generation of talent and forfeits the spillovers that make small economies rich. For a mid-sized open economy inside Horizon Europe, foresight is the discipline that earns the right to say no to thirty me-too priorities and concentrate.

### 15–20y

Horizon — capabilities, instruments & people that must exist now

### Frontier

Signal — patent-citation lead-lag & talent-flow clustering before a field is fashionable

### Roadmapping agent

Maintains a living capability-to-mission dependency graph, re-scoring feasibility as evidence lands

### FIRST MOVE

Stand up the scanning agent and publish a baseline frontier map; run one structured Delphi to rank candidate missions; deliver a defund-and-concentrate proposal into the next budget round. **Owner:** Research Priorities Foresight Unit inside the national R&D council.

## 10 Trade, Supply-Chain & Geoeconomics

Mapping chokepoints before they close — the single mines, fabs, ports and protocols on which prosperity silently rests, and how decoupling and friend-shoring will redraw them.

### WHY IT RANKS

For a small, deeply open, export-led economy braided into German automotive and European value chains, trade is the circulatory system. The lead time to qualify a second supplier or reshore a process runs to years; the lead time of a shock runs to days. Foresight closes that gap before the bill arrives.

### 1–3y

Horizon — which strategic inputs have no qualified alternative supplier

### Lead-time creep

Signal — bottleneck inventories & freight anomalies as early tremors

### Red-team agent

Plays the adversary, designing the coercion campaign that would hurt most

**Method** Scenario worlds: managed decoupling · bloc consolidation · fragmentation

**KPI** Strategic inputs with a qualified second source

**Owner** Geoeconomic Foresight Cell · industry + NSA

**First move** Map two flagship value chains to component level

## 11 Digital & Data-Infrastructure

Provisioning the national digital backbone ahead of demand — treating compute and data capacity the way a serious state treats grid and water: a strategic stock to be planned, not a service procured when the lights flicker.

### WHY IT RANKS

It is the substrate beneath innovation, public-service quality, financial supervision and defence. Lead times are long and dependencies sovereign: GPU supply is rationed globally, data-centre power collides with grid and climate limits, and cloud dependence becomes geopolitical exposure. Digital infrastructure is a capacity-planning problem on the scale of energy policy, not an IT budget line.

### 2–4y

Horizon — where compute, bandwidth & skilled-ops demand outruns supply

### Grid queues

Signal — data-centre power & cooling as the binding constraint

### Wind-tunnel agent

Runs the digital-government roadmap through demand futures, flags capacity it assumes but will not exist

### FIRST MOVE

Build the national compute supply-demand model and publish the first capacity outlook; run one backcasting exercise to a 2032 capacity target. **Owner:** Digital Infrastructure Foresight function inside the digital ministry, nested in EU sovereign-cloud and AI-factory initiatives.

## 12 Financial-System & Systemic-Risk

Seeing the next crisis form — watching the system the way an epidemiologist watches a population: not whether each institution is sound, but where fragility is quietly accumulating and how a small failure cascades.

### WHY IT RANKS

Financial stability is the precondition for everything else: a banking or sovereign-debt crisis vaporises years of growth faster than any other shock. The damage is non-linear and the warning window is the thing always missing. Stress tests ask "can the system survive the last crisis?"; foresight asks "what is the next one, and is anyone mandated to look?"

### 1–2y

Horizon — where leverage builds outside the regulated banking core

### Credit gap

Signal — credit growth vs. trend, the most reliable banking-crisis predictor

### Early-warning agent

Watches the live vulnerability map against historical crisis signatures, fires graded alerts



*Financial crises are never new in substance and always new in costume.*

ENSI FORESIGHT DIVISION · AREA 12 · OWNER: SYSTEMIC-RISK FORESIGHT UNIT, CENTRAL BANK

## 13 Infrastructure & Capital-Planning

Building 50-year assets for a future we can only imagine — reframing the capital pipeline as a portfolio of options under deep uncertainty: not "what will demand be in 2055?" but "which assets perform across the futures we cannot rule out?"

### WHY IT RANKS

Infrastructure errors are uniquely unforgiving: a stranded grid corridor or an under-sized lock chamber is a fifty-year liability cast in concrete, with sunk cost defending it against correction. Decarbonisation is rewiring demand, climate is rewriting the load envelope, and reshoring is redrawing the European freight map — capital is being committed now against the most uncertain fifty years in living memory.

### 30–50y

Horizon — asset life, with 10-year adaptation checkpoints

### Utilisation gap

Signal — realised-vs-forecast demand on opened assets, the cheapest lesson

### Simulation agent

Re-runs RDM & adaptive-pathway models nightly; flags assets slipping robust→fragile

### FIRST MOVE

Stand up the DMDU cell, retrofit a robustness gate into appraisal guidance, run two pilot reappraisals — a grid corridor and a water asset. **Owner:** finance + transport ministries, mirroring the UK IPA × HM Treasury pairing.

## 14 Critical Minerals, Materials & Resource-Security

Securing the inputs the future economy runs on — lithium, rare earths, gallium, high-purity silicon, copper — whose supply is concentrated, politicised and slow to scale. The supply-side insurance policy for every industrial ambition.

### WHY IT RANKS

Materials are a chokepoint multiplier: a single export restriction can stall an entire downstream industry that took a decade to build. The 2023 gallium and germanium controls showed how a refining monopoly converts into geopolitical power overnight. Opening a mine takes 10–15 years; a supply shock lands in weeks — foresight is the only instrument operating on the timescale of the cure.

### 2–5y

Horizon — acute disruption & stockpile adequacy

### Refining HHI

Signal — who controls processing, not just reserves

### Scanning agent

Monitors export-control filings, customs flows & price feeds across languages in near-real-time

**KPI** Stockpile-adequacy days for designated critical inputs

**Owner** Resource-security unit · industry + defence + EU CRMA

**First move** Red-team battery-grade lithium / rare-earth magnets

## 15 SME, Entrepreneurship & Firm-Dynamics

Anticipating the future of the firm base — where the productive capacity and jobs of 2035 actually form, which kinds of firm, in which sectors, at what scale — and shaping the environment to meet that emergence, not subsidise its predecessor.

### WHY IT RANKS

The firm base is where every other lever cashes out into productivity and employment — and OECD business dynamism has fallen for two decades: fewer high-growth entrants, more zombie incumbents. Czechia's maturing manufacturing base, thin scale-up layer and retiring founder-owners make this a live strategic question. AI now threatens to reshape firm size faster than any institution can react.

### 3–10y

Horizon — dynamism trends & the AI-driven anatomy-of-the-firm shift

### HGF share

Signal — high-growth-firm share & frontier-vs-laggard productivity gap

### Simulation agent

Agent-based firm-population models test how a tax or scheme reshapes entry, scaling & exit

### FIRST MOVE

Build the secure firm-level microdata pipeline and a baseline dynamism dashboard; pilot a founder-succession early-warning-and-support scheme for the retirement wave. **Owner:** industry ministry + SME agency, fused to the statistical office.

## 16 Foreign Investment & Strategic-Sector

Getting ahead of where global capital moves next — positioning to attract the strategic sectors of the future while screening the inbound flows that threaten sovereignty. National positioning in a contest already underway.

### WHY IT RANKS

Investment-attraction is a winner-take-much contest with durable lock-in: a gigafactory or AI datacentre cluster anchors suppliers, skills and follow-on capital for decades, and there are only so many to win. Reacting to an investment wave is losing it — the site-selection, incentives and skills pipeline must be ready before the capital looks. The EU FDI-screening framework exists because some inbound capital is a vector for strategic dependence.

### 3–7y

Horizon — relocation & investment waves as supply chains reconfigure

### Site shortlists

Signal — greenfield FDI scouting & competitor incentive packages

### Red-team agent

Plays rival investment-promotion agencies, exposing how a competitor out-bids for a target

**KPI** Win-rate on targeted strategic investments vs. rival jurisdictions

**Owner** FDI-foresight cell · promotion agency + inter-ministerial screening

**First move** War-game the next relocation wave; codify anticipatory screening

TIER III · AREAS 17-24

# Society, Resilience & Quality of Life

Where foresight touches how people actually live — health, learning, ageing, migration, cohesion, place, housing and wellbeing — and where anticipation is the difference between a society that adapts and one that fractures.

17 Health & Pandemic

18 Education & Lifelong-Learning

19 Ageing & Longevity

20 Migration & Integration

21 Cohesion & Trust

22 Regional & Place-Based

23 Housing & Urban-Systems

24 Wellbeing & Beyond-GDP

## 17 Health System & Pandemic

Designing a health system for the demand of 2040, not 2010 — treating demography, epidemiology, biosecurity and workforce as one coupled system, and asking not "is it efficient now?" but "is it solvent against 2040?"

### WHY IT RANKS

Health is a growth engine, a stability anchor and the most visceral component of quality of life — and its failure modes are non-linear. A pandemic erases years of GDP in months; an under-projected workforce gap closes hospital wards a decade after the missed recruitment decision. WHO, ECDC and HERA exist because reactive health policy is ruinously expensive.

### 0–25y

Horizon — surge capacity now → AI-and-AMR-shaped system later

### Wastewater

Signal — genomic surveillance flagging anomalies against baseline

### Demand-projection agent

Keeps the workforce-and-capacity model live, surfacing the gap years before it bites

### FIRST MOVE

Stand up genomic-and-wastewater scanning on existing surveillance; publish the workforce-gap projection to 2040; run a pandemic wind-tunnelling exercise across three shock scenarios. **Owner:** Health Foresight Unit, ministry of health, wired to ECDC/HERA.

## 18 Education & Lifelong-Learning System

Teaching for a world that doesn't exist yet — a child entering school this year will work into the 2070s. Teach forward: anticipate the skills, the shape of work, and the rhythm of learning across a whole life.

### WHY IT RANKS

Skills are the binding constraint on everything else in this report — you cannot run an energy transition, a digital economy or a modern health system without the people to staff them, and people take fifteen to twenty years to grow. Get it wrong and you import the skills, watch your young emigrate, or stall the transitions you committed to. The recurring question: are we teaching content AI will commoditise, or capabilities that compound?

### 3–10y

Horizon — how automation reshapes which tasks humans do

### Adult-learning rate

Signal — best proxy for whether lifelong learning is real or rhetorical

### Curriculum red-team agent

Stress-tests proposed standards: which competencies are robust, which bet on one future?

**KPI** Adult-learning participation, on a path to the Nordic frontier

**Owner** Education & Skills Foresight Unit · education + labour

**First move** Live emerging-skills map; futures-literacy pilot for teachers

## 19 Ageing, Care & the Longevity Economy

Turning a demographic burden into an economic frontier — making care, pension and fiscal systems solvent against the curve while capturing the longevity-economy market that older, healthier, wealthier populations open up.

### WHY IT RANKS

Ageing is the master demographic driver beneath health, care, pensions, labour supply and public finance — and unusually certain: the people old in 2040 are already alive and counted. Left reactive, it is a slow fiscal landslide; handled with foresight, the same demographics become a stability story and a growth story. For Czechia, with one of the EU's steeper ageing curves, this is the central fiscal question of the next two decades.

### 10–25y

Horizon — life, work & care at a higher median age

### HLE gap

Signal — gap between living longer and living well is the whole game

### Simulation agent

Live demographic-fiscal model: pension solvency & care demand as a running dashboard

### FIRST MOVE

Build the integrated demographic-fiscal-care model and publish an honest solvency projection to 2050; commission a longevity-economy opportunity map turning the curve into industrial strategy. **Owner:** Ageing & Longevity Foresight Unit, spanning finance, health and labour.

## 20 Migration & Integration

Anticipating flows and making integration work — reading the drivers that move people before flows crest, then ensuring arrivals find their way into the labour market and social fabric. Arrival and integration as one system to be designed.

### WHY IT RANKS

Migration is a sharp instrument that cuts toward growth or instability depending entirely on management. For ageing, shrinking European societies, managed inward migration is among the few levers that ease workforce and fiscal pressure. The long-run outcome is decided not at the border but in the decade after — and integration outcomes shape social trust for a generation.

### 3–10y

Horizon — do integration systems deliver employment, or warehouse people?

### Driver indices

Signal — conflict, climate-stress & economic-collapse data upstream of flows

### Integration model agent

Agent-based labour-matching: where does this cohort end up in five years?

**KPI** Migrant-native employment gap, and whether it is closing

**Owner** Migration Foresight Unit · interior + labour + local government

**First move** Driver early-warning agent; integration-outcomes dashboard

## 21 Social Cohesion, Inequality & Trust

Spotting the fracture lines before they break — treating the fabric of society as a forecastable system, watching where trust thins and modelling how it propagates while intervention is still cheap and legitimate.

### WHY IT RANKS

Cohesion is the precondition for everything else: growth strategies, green transitions and demographic adjustments all demand that citizens accept disruption today for benefit tomorrow — and acceptance runs entirely on trust. Where trust erodes, even good policy is read as elite capture and reforms stall. A country that anticipates where cohesion thins gains the social licence to actually govern.

### 2–4y

Horizon — where trust & perceived fairness deteriorate fastest

### Perception gap

Signal — gap between objective conditions and perceived fairness moves people

### Opinion-dynamics agent

Live agent-based model stress-tests how a shock propagates through specific communities

### FIRST MOVE

Stand up a shared cohesion data spine; commission a CLA-driven scenario set; publish an annual "State of Social Cohesion" report that forces the slow variable onto the cabinet table before it breaks. **Owner:** Social Cohesion unit, centre of government, insulated from party spin.

## 22 Regional & Place-Based Development

Keeping the whole country in the future — refusing the comfort of the national average, surfacing regional divergence before it hardens into permanent decline, and investing in opportunity geographies while the window is open.

### WHY IT RANKS

The geography of opportunity has become the defining political and economic cleavage of the age — the "places that don't matter" striking back through the ballot box. For Czechia, the Prague-versus-Ústí/Karlovy Vary/Moravian-Silesian cleavage is acute, and EU Cohesion Policy moves enormous sums blind to long-horizon trajectories. Divergence destroys productivity, corrodes cohesion and degrades quality of life at once.

### 3–5y

Horizon — which places sit at a tipping point well-timed investment can change

### 25–40 migration

Signal — internal flows of the prime-age band, ahead of census confirmation

### Wind-tunnel agent

Tests a Cohesion-fund allocation against every scenario, exposing waste on managed decline

**KPI** Regional convergence — narrowing GDP & net-migration gaps

**Owner** Regional-development ministry + statistics office + mayors

**First move** Build the territorial data twin; publish a regional futures atlas

## 23 Housing & Urban-Systems

Planning the city of 2050 while it is still affordable to plan — cities are slow machines built from fast decisions; a zoning rule or height limit set today locks in form, cost and liveability for fifty years.

### WHY IT RANKS

Housing is where the macroeconomy meets the kitchen table: unaffordable housing suppresses fertility, traps labour, transfers wealth violently across generations, and drives the cohesion and regional-divergence crises directly. Prague's affordability is among Europe's worst relative to incomes, and supply responds on a five-to-ten-year lag — by the time a problem is obvious it is a generation too late to fix cheaply.

### 15–25y

Horizon — the urban form ageing, climate & remote work require

### Price-to-income

Signal — house-price & rent ratios vs. the pipeline-to-demand gap

### Land-use sim agent

Coupled land-use/transport/market models project affordability & emissions per scenario

### FIRST MOVE

Build the city-region digital twin for Prague and one secondary city; backcast to a 2050 affordable-city vision with municipal leaders; wind-tunnel current zoning. **Owner:** regional-development ministry + major-city planning authorities.

## 24 Public Health, Wellbeing & Beyond-GDP

Steering by quality of life, not just output — anticipating the trajectory of healthy-life-expectancy, mental health and lived quality of life, and building the dashboards that let governments optimise for what citizens experience.

### WHY IT RANKS

Wellbeing is both end and means: healthy, secure, mentally well populations are more productive, cohesive and resilient — quality of life is a precondition of sustainable growth, not a luxury bought after it. The deterioration is underway: rising youth mental-ill-health, stalling healthy-life-expectancy, an epidemic of loneliness, and a widening gap between rising output and flat life-satisfaction.

### 10–15y

Horizon — how ageing & the digital transformation reshape healthy-life-expectancy

### GDP–wellbeing gap

Signal — the widening divergence is itself the key early warning

### Wellbeing observatory agent

Tracks quality of life in near-real-time, not lagging annual reports

**KPI** Healthy-life-expectancy – raise it, close the gap to total LE

**Owner** Health ministry + centre of government + statistics office

**First move** Official beyond-GDP dashboard beside the fiscal numbers

TIER IV · AREAS 25-32

# Environment, Security & the Enabling Layer

The planetary and security boundaries the economy runs inside, the anticipatory governance that keeps law ahead of technology, and — load-bearing for the whole report — note [31](#) (the data backbone) and [32](#) (futures literacy), the substrate every other area executes on.

[25](#) Climate Adaptation

[26](#) Food, Water & Agriculture

[27](#) Biodiversity & Natural Capital

[28](#) Defence & Security

[29](#) Cyber & Information Integrity

[30](#) Anticipatory Regulation

[31](#) Foresight Data Architecture

[32](#) Futures Literacy & Culture

## 25 Climate Adaptation & Resilience

Pricing in a future climate that is already locked in — sizing roads, grids, hospitals and fiscal buffers for the climate of the 2050s, not the 1990s in which they were planned.

### WHY IT RANKS

Adaptation is the resilience floor under every growth ambition: a flooded substation or a low-flow river starving a plant of cooling water erases years of output in a week. For a landlocked, river-dependent industrial economy — Elbe and Morava flooding, Moravian drought, grid heat stress — adaptation *is* macro-fiscal policy. Physical risk transmits into transition risk, insurance withdrawal and sovereign borrowing cost.

### 0–40y

Horizon — from mispriced assets today to managed retreat under 3°C-plus worlds

### Signal

Insurer premium & withdrawal patterns — the spread of "uninsurable" postcodes

### Agent

Simulation agents translate a warming path into asset-level damage & fiscal exposure

#### FIRST MOVE · 12 MONTHS

Stand up an asset-level physical-risk model and a live national exposure map; wind-tunnel the sovereign budget against three warming scenarios; publish a ranked least-regret adaptation pipeline tied to the capital plan. **Owner:** Climate Resilience Unit (environment + finance), wired to the central bank's financial-stability function.

## 26 Food, Water & Agricultural Systems

Securing the non-substitutable basics under stress — turning "we are mostly self-sufficient" complacency into a scenario-tested map of where the real fragilities sit.

### WHY IT RANKS

Food-price spikes are among history's most reliable precursors of unrest, and even import-secure economies are exposed through global commodity, fertiliser and energy markets — the 2022 shock chained them together fast. Drought-prone, CAP-dependent Central Europe, with falling groundwater and stressed soils, faces this as a direct driver of rural stability and food-price inflation that monetary policy cannot easily tame.

### 0–30y

Horizon — from the next price shock's origin to which crops & regions stay viable

### Signal

Export-restriction announcements + grain-reserve levels (FAO, JRC MARS bulletins)

### Agent

Simulation agents trace a fertiliser spike or drought to domestic shelves & farm balance sheets

#### FIRST MOVE · 12 MONTHS

Build a live food-and-water dashboard fusing crop, water, price and trade signals; run a compound supply-shock simulation (drought + export ban) against reserves; publish a costed agricultural-transition pathway tied to CAP funding. **Owner:** Food & Water Security Cell (agriculture + water authorities + environment), linked to the central bank for the inflation channel.

## 27 Environment, Biodiversity & Natural Capital

Managing the productive assets that don't appear on the balance sheet — forests, soils, pollinators and wetlands that are silently liquidated until the services they provide fail.

### WHY IT RANKS

A very large share of economic activity depends moderately or highly on nature, and the danger is non-linearity: an ecosystem degrades gradually, then collapses, wiping out services assumed free and permanent. The bark-beetle devastation of spruce monocultures was a vivid, expensive lesson — a managed ecosystem crossing a threshold and imposing sudden costs on the timber economy and landscape, often irreversible on policy timescales.

### 0–50y

Horizon — from species nearing thresholds to a nature-positive economy

### Signal

Pest outbreaks, fishery collapse & invasive-species spread marking thresholds approached

### Agent

Scanning agents fuse satellite, sensor & field data to detect degradation before periodic surveys



*Reframe environmental policy from a compliance chore into asset management of the stock the economy quietly depends on — a stock that, once past a threshold, cannot be rebuilt with money.*

METHOD ANCHOR · CAUSAL-LAYERED ANALYSIS + NATURAL-CAPITAL ACCOUNTING (UN SEEA, TNFD)

## 28 Defence & National Security

Reading the strategic environment before it turns — matching slow-to-build force design, posture and warning to the war that might come, not the last one fought.

### WHY IT RANKS

Security is the precondition for everything else, and being wrong here is least recoverable: forces take a decade to build into a threat environment that may already have transformed. The post-2022 environment — Russian aggression, rapid rearmament, the drone-and-EW revolution, a contested NATO eastern flank — gives a frontline-adjacent member a direct stake in getting force design and warning right. Long lead times make foresight structurally necessary, not optional.

### 0–30y

Horizon — from warning indicators of aggression to the enduring future-force base

### Signal

Grey-zone activity — cyber intrusion, GPS jamming, undersea-cable & energy-infra incidents

### Agent

Wargaming agents stress-test force designs and expose escalation paths planners can't enumerate

## 29 Cyber, Hybrid-Threat & Information Integrity

Defending the information space and the democratic core — staying ahead of a threat surface that mutates monthly and runs at machine speed.

### WHY IT RANKS

This is a precondition, not a sector: a compromised election, a poisoned public conversation or a paralysed hospital network corrodes the trust every other policy depends on. Because attacker innovation outpaces defender procurement cycles, only an anticipatory posture closes the gap — reactive cyber-defence guarantees you are always one campaign behind. The question is whether democratic resilience compounds faster than adversary capability.

### 0–15y

Horizon — from AI-scaled disinformation to whether the open commons survives

### Signal

Cost-per-thousand of credible synthetic content; new offensive AI tooling on state/criminal forums

### Agent

Detection agents flag coordinated inauthentic behaviour & synthetic-media signatures near-real-time

### FIRST MOVE · 12 MONTHS

Stand up the always-on scanning-and-simulation platform; run a full hybrid-campaign wargame against the next election; publish a recurring national information-threat foresight brief. **Owner:** Information-Integrity & Cyber-Foresight cell (national cyber authority — NÚKIB — + strategic-communications), wired to ENISA, EEAS and the EU Rapid Alert System.

## 30 Anticipatory Regulation & Governance

Regulating emerging technology ahead of the harm — flipping the default from legislating after the crash to shaping rules, sandboxes and institutions before harm is locked in.

### WHY IT RANKS

Regulation is the interface between every emerging technology and the real economy; a slow or brittle interface throttles both safety and growth. Reactive rule-making delivers the worst of both worlds — harm runs ahead unchecked, then a panicked overcorrection chills investment. Inside the EU's regulatory gravity (AI Act, GDPR, DSA/DMA), anticipatory governance is how a smaller state punches above its weight: the jurisdiction where adaptive rules make responsible innovation faster, not slower.

### 0–15y

Horizon — from near-term regulatory blind spots to a self-revising regulatory state

### Signal

The gap metric — lag between a technology reaching market and a coherent rule existing

### Agent

Simulation agents wind-tunnel proposed rules to predict compliance behaviour & loopholes pre-enactment

# 31 The Foresight Data Architecture & Decision-Intelligence Platform

The shared backbone that powers every other area — a national data architecture and decision-intelligence layer that turns raw statistics, sensors and registers into integrated, near-real-time, decision-ready signals.

## WHY IT RANKS – THE ENABLING LAYER

Nothing else works without it. A scenario is only as credible as the indicators feeding it; an early-warning system that cannot see across health, energy, mobility and finance in one view is blind precisely where modern crises form. The World Bank's *Data for Better Lives* made the core case: data is a public good whose value compounds when integrated and reused.

For the Czech Republic, with a capable national statistical office (the **Czech Statistical Office, ČSÚ**) and the **UNECE HLG-MOS** modernisation agenda already in play, the backbone is within reach — but only if treated as strategic infrastructure with sustained investment, not a perennial IT afterthought.

## METHODS THAT FIT

A *data-mesh* approach lets each ministry own its domain data while exposing it as a governed, interoperable product — federation without a single brittle warehouse. *Knowledge graphs* provide the semantic layer linking people, places, events and indicators across silos. *Nowcasting* converts high-frequency sources — satellite, mobility, payments, sensors, web-scraped prices — into live reads of the economy, health and environment. The standard scanning, scenario and cross-impact methods all draw from this common, trusted substrate.

### 0–15y

Horizon — from unblocking decisions made blind today to a fully agentic, privacy-preserving platform

### Signal

Decision latency — how fast a signal becomes an actionable indicator on a decision-maker's screen

### Agent

The agentic engine *is* the platform — ingestion, linking, nowcasting, query & privacy-preserving agents



*The statistical office evolves from periodic publisher into the operator of an always-on, trustworthy decision-intelligence platform — the engine room of the foresight state.*

SHARED SUBSTRATE FOR ALL 31 OTHER AREAS · REFERENCES **ČSÚ** & **UNECE HLG-MOS** MODERNISATION

## FIRST MOVE · 12 MONTHS

Stand up the data-mesh and knowledge-graph foundation for three priority domains; launch a nowcasting pilot for the economy and one public-service area; codify the data-governance and privacy framework that lets the backbone scale with public trust intact. **Owner:** the national statistical office (ČSÚ) with the central digital-government authority and a cross-government data-governance board; connects to EU data spaces and Eurostat.

## 32 Futures Literacy, Culture & Capability-Building

The human operating system for a foresightful state — the skills, habits and culture that turn the other thirty-one areas from reports into reflexes.

### WHY IT RANKS – CULTURE EATS METHOD

Foresight fails most often not because the analysis was wrong but because no one in the room knew how to use it. A state can buy the best scanning platform and commission brilliant scenarios, yet if officials default to short-term, reactive, single-future thinking, the investment is wasted. UNESCO made futures literacy a recognised competence precisely because the bottleneck is human, not technical.

The differentiator is whether anticipation is a specialist unit's job or a shared way of working. The countries that win are not those with the most foresight reports, but those where foresight has become an ordinary, embedded capability of the people who govern.

### METHODS THAT FIT

UNESCO-style *Futures Literacy Laboratories* let people rigorously imagine and use multiple futures; reframing exercises surface and challenge the implicit assumptions behind today's decisions. Participatory scenario workshops and the *Futures Wheel* build the muscle by doing, not lecturing. Causal Layered Analysis trains officials to see beneath events to the deeper narratives — paired with communities of practice and "learning by intervening" on real policy problems, not abstract rehearsal.

### 0–15y

Horizon — from the minimum capability that unblocks decisions to a self-sustaining culture

### Signal

The decisive ratio — foresight that informs a real decision vs. foresight that merely gets published

### Agent

Tutoring, facilitation & coaching agents lower the cost of being foresightful — judgement stays human



*Coaching agents prompt the anticipatory question inside real decision workflows — "what would have to be true in ten years?" — until anticipation becomes the path of least resistance.*

CAPABILITY NETWORKS · UNESCO · OECD FORESIGHT NETWORK · SOIF · POLICY HORIZONS CANADA

### FIRST MOVE · 12 MONTHS

Launch a futures-literacy curriculum and a cohort of trained foresight facilitators seeded across ministries; embed a lightweight anticipatory step into existing planning and budgeting cycles; stand up a cross-government community of practice so the capability compounds rather than depending on a single champion. **Owner:** the public-administration academy with the centre-of-government foresight function, wired into HR and leadership development.

CASES · THE MODEL END-TO-END

# Foresight states, working

Four governments did not get better forecasts than everyone else — forecasts are usually wrong. They built the *capability*: a standing engine that scans, models, stress-tests, and routes what it learns into the decisions being made anyway. Each proves a different transferable lesson about how anticipation is wired into power.

01 · SINGAPORE

## A whole-of-government futures engine in the PMO

The Prime Minister's Office runs the **Centre for Strategic Futures (CSF)**, paired with the Strategic Futures Network of permanent secretaries. Its method lineage — **RAHS** (Risk Assessment & Horizon Scanning) and **Scenario Planning Plus** — turns scanning into the national scenario set every ministry plans against.

**Routes via** The PMO seat plus the permanent-secretary network make foresight a required input to the budget and strategy cycle — not a courtesy memo.

**Lesson** Put the engine at the centre, with a standing convening mandate, so it cannot be defunded when a champion moves on.

SINCE

2009 CSF stood up after the RAHS programme · scenario practice traces to the 1990s Shell-influenced Scenario Planning Office · *Foresight* journal now in its 15th-anniversary cycle.

### PMO

Seat at the cabinet centre

### Whole-gov

One shared scenario set

### 15+ yrs

Continuous practice

02 · FINLAND

## Foresight wired into the legislature itself

Finland is unique in giving foresight a **parliamentary** home: the **Committee for the Future**, a standing committee of MPs that debates the government's quadrennial **Report on the Future**. The innovation fund **Sitra** supplies the megatrends and futures capability that feed it.

**Routes via** The government must publish a long-horizon Report on the Future each term; Parliament must formally respond — anticipation becomes a constitutional habit, not an option.

**Lesson** Anchor foresight on *both* sides of power. A legislative hook outlasts any single administration.

SINCE

1993 Committee for the Future established · 1988 Sitra's futures work matures · Government Report on the Future now a fixed once-per-term obligation.

### Statute

Parliamentary committee

### Per term

Mandatory future report

### Sitra

Standing megatrends

03 · CANADA

## A dedicated, disciplined scanning organisation

**Policy Horizons Canada** is the federal government's own foresight organisation, producing the flagship annual **Disruptions on the Horizon** — a structured, ranked scan of plausible

04 · UNITED ARAB EMIRATES

## A sovereign futures agency that ships

The **Dubai Future Foundation** treats the future as something to be prototyped, not just studied — convening ministries, running the Museum of the Future, and pushing foresight straight into pilots, regulation-light testbeds and sovereign

CASES · WHAT THE FOUR PROVE TOGETHER

# Four routes into power, one pattern

Read side by side, the cases are not four models to choose between — they are four answers to the same question: *how does anticipation get a guaranteed route into the decision?* Singapore puts it at the executive centre; Finland anchors it in the legislature; Canada makes it a disciplined public craft; the UAE fuses it with delivery. Each solves the failure mode the others are exposed to.

STATE	WHAT THEY BUILT	THE ROUTE INTO DECISIONS	TRANSFERABLE LESSON
Singapore	CSF in the PMO; RAHS + Scenario Planning Plus; Strategic Futures Network of perm-secs.	Cabinet-centre seat; shared national scenario set feeds budget & strategy.	Build the engine at the <b>centre</b> — one scan for the whole of government.
Finland	Parliamentary Committee for the Future; Government Report on the Future; Sitra megatrends.	Statutory once-per-term future report; Parliament must respond.	A <b>legislative hook</b> makes foresight survive elections.
Canada	Policy Horizons Canada; <i>Disruptions on the Horizon</i> (35 disruptions, 2024).	Repeatable annual product; open methods + training across departments.	<b>Cadence and method</b> beat one-off brilliance.
UAE	Dubai Future Foundation; Museum of the Future; testbeds and moonshots.	Leadership sponsorship; scenarios convert into pilots and regulation-light testbeds.	Pair foresight with <b>delivery</b> — prototype the future.

THE COMMON THREAD

None of the four bought a forecast. Each bought **option value**: the ability to see a shift early, to have already rehearsed the response, and to act while action is still cheap. The institutional design differs; the engine — *scan* → *model* → *stress-test* → *route* — is the same in every one.

WHAT A COUNTRY LIKE CZECHIA CAN BORROW

Take Singapore's **centre-of-government seat**, Finland's **statutory hook**, Canada's **repeatable cadence**, and the UAE's **bias to build** — and skip the failure modes each was built to avoid. The model is portable; the agentic layer makes it cheap.



*The countries that built foresight did not get better forecasts than everyone else. What they got was the ability to act while action was still cheap.*

THE ACTIONABLE PAYOFF

# The 12-month build roadmap

The order is the strategy. A country starting from zero stands up the **minimum viable foresight state** first — the centre-of-government engine, the national early-warning system, the data backbone, and futures literacy — then extends along the macro spine. Build the engine once; every later domain is a low-cost extension of a system that already works.

**START MONDAY**

You do not need a law, a budget line, or a new agency to begin. Name a **Chief Foresight Officer** reporting to the cabinet secretary, convene a lean directorate of eight to twelve, and publish a **first national scenario set** with a standing signals registry behind it. Wind-tunnel the next budget against it. Everything else compounds on that first, cheap move.

FIRST MOVE	OWNER	HORIZON	METRIC TO WATCH
<b>1 · Stand up the foresight engine</b> — a lean directorate at the cabinet centre; publish the first national scenario set; wind-tunnel the next budget.	Chief Foresight Officer → cabinet secretary <b>Area 1</b>	Q1	Share of major decisions wind-tunnelled; signal-to-response lead time.
<b>8 · Stand up national early-warning</b> — always-on cross-hazard scanning; rebuild the risk register as a living, cascade-based system.	Resilience / national-security secretariat (PMO) <b>Area 8</b>	Q1-Q2	Warning lead time on materialised risks; preparedness coverage of priority risks.
<b>31 · Lay the data backbone</b> — data-mesh + knowledge graph for three priority domains; one nowcasting pilot; codify the privacy framework.	National statistical office + digital-gov authority <b>Area 31</b>	Q1-Q3	Decision latency (event-to-indicator); interoperable data products published.
<b>32 · Build futures literacy</b> — a curriculum + a cohort of trained facilitators seeded across ministries; a cross-government community of practice.	Civil-service academy + centre-of-gov foresight <b>Area 32</b>	Q2-Q3	Share of officials with working futures literacy; foresight-to-decision ratio.
<b>2 · Extend to the fiscal spine</b> — replace single-baseline budgeting with a three-scenario stress test; build the contingent-liability register.	Finance ministry + independent fiscal council <b>Area 2</b>	Q3-Q4	Contingent liabilities priced; medium-term forecast error narrowed.
<b>3 · Extend to growth &amp; industry</b> — build the economic-complexity dashboard; shortlist reachable high-value rungs; deliver a diversification map.	Economy ministry strategy fn + development bank <b>Area 3</b>	Q4	Complexity index trajectory; relatedness density around target activities.
<b>5 · Extend to labour &amp; skills</b> — stand up the live vacancy-and-exposure dashboard; publish a first three-horizon skills outlook.	Labour + education ministries; PES <b>Area 5</b>	Q4	Skills-mismatch index; share of training places aligned to forecast demand.

**THE SEQUENCING RULE**

Attempt the society and environment tiers *before* the engine and you get isolated reports nobody acts on. Build the engine first — Areas **1, 8, 31, 32** — and each new domain is a cheap extension of a system that already works.

THE FIRST YEAR, PHASED

# The minimum viable foresight state

Four moves make a foresight state that functions: an engine at the centre, eyes on the horizon, a spine of data, and people who can think in futures. Stand these up in the first year and the macro spine plugs straight in. The phasing below turns the roadmap into a calendar.

- 1 **Quarter 1 — Stand up the engine and the watch.** Name the Chief Foresight Officer; recruit the directorate; publish the first national scenario set and the signals registry; switch on always-on cross-hazard scanning. The state can now see, and has one shared picture of the future.
- 2 **Quarter 2 — Wire it into decisions and people.** Embed foresight liaisons in finance, economy and interior; wind-tunnel the first budget input; launch the futures-literacy curriculum and the first facilitator cohort. Anticipation starts becoming a habit, not an event.
- 3 **Quarter 3–4 — Lay the backbone and extend the spine.** Stand up the data-mesh and knowledge graph for three domains; run the first nowcasting pilot; extend foresight into the fiscal, growth and skills functions with their own dashboards. The engine is now load-bearing.

4

Foundational areas (1 · 8 · 31 · 32) = the minimum viable foresight state

8–12

People in the founding directorate

12 mo

From zero to a functioning engine

0

New laws required to start

PHASE 1 · Q1

See

Engine + scenario set + signals registry + cross-hazard watch. The state stops being surprised by the predictable.

PHASE 2 · Q2

Route

Liaisons in ministries; first wind-tunnelled budget; futures literacy seeded. Foresight reaches the decision.

PHASE 3 · Q3–4

Compound

Data backbone live; macro spine plugged in. Each new domain is now a cheap extension.

THE WAGER, RESTATED

Anticipation is the cheapest investment a state can make: a modest standing cost to see a shift early, set against the large, lumpy cost of being surprised by it late. Compounded across 32 domains and across decades, that gap is the difference between a country carried by the future and one that authors it.

CREDIBILITY THROUGH HONESTY

# Risks, counter-arguments & steelman

A foresight capability that cannot survive its own critics does not deserve a budget. So here are the five objections raised most often — each stated at its strongest, the way a sceptical finance minister would put it — and then answered. The honest answer to most of them is not "you're wrong" but "you've described foresight done badly, which is exactly what this design prevents."

THE OBJECTION (STEELMANNED)	THE ANSWER
"Foresight is just guessing." The future is unknowable; dressing up guesses as scenarios lends false authority to what is, at bottom, speculation.	Correct — and beside the point. Foresight does not claim to predict; it claims to <i>prepare</i> . Its product is not a forecast but <b>option value</b> : the rehearsed response, the early move, the assumption tested before money is committed. You do not need to know which crisis comes to be less surprised by all of them.
"Agents hallucinate, and automation bias makes it worse." An AI fleet that confidently invents signals, scaled across government, is more dangerous than no foresight at all.	Which is why the design is explicit: agents <b>widen the aperture; humans keep the accountability</b> . Agents scan, generate and red-team; a human directorate curates, judges and owns what goes up. Every claim is traceable to source, and red-team agents are pointed at the system itself. Discipline is a design choice, and we made it.
"It'll be captured by short-termism — or become theatre." Foresight units get colonised by the politics of the moment, or survive as ritual workshops that change nothing.	The single most common failure mode — and the reason for the <b>statutory hook</b> (Finland) and the <b>centre-of-government seat</b> (Singapore). The KPI that matters is the ratio of foresight that <i>changes a decision</i> to foresight that merely gets published. Measure that, and theatre has nowhere to hide.
"We can't afford it." Fiscal slack is tight; a new unit is a cost we can defer until conditions improve.	The asymmetry runs the other way. A founding directorate is <b>eight to twelve people</b> ; the cost of one un-anticipated crisis — a skills gap that becomes mass unemployment, a risk that becomes an emergency, an asset that strands — dwarfs it. For a mid-sized economy with little fiscal slack, the asymmetry is <i>sharper</i> , not weaker.
"Scenarios never come true." We built scenarios last cycle and none described what actually happened — so why build more?	A scenario that "comes true" was a forecast, and a lucky one. Scenarios are not bets on one future; they are <b>stress tests</b> for many. The pandemic was on every risk register and the energy shock in every scenario set — the failure was never imagination, it was <i>routing</i> . This design fixes the routing.

THE DEEPER RISKS · AND HOW THE DESIGN ANSWERS THEM

# Why foresight units die — and the guardrails

Behind the five objections sits one real danger: not that foresight is wrong, but that it is *built badly*. Foresight units die from a small, repeating set of causes. The guardrails below are the response — each a structural answer, not a slogan.

## The four ways foresight dies

- ▶ **Orphaned reports.** Brilliant scenarios with no route into the budget or the legislative cycle — read once, shelved forever.
- ▶ **No mandate.** A unit that can convene no one, so its work is a courtesy ministries are free to ignore.
- ▶ **No data.** Scenarios built on stale, ministry-by-ministry data that cannot be combined when a decision is due.
- ▶ **No people.** A culture that defaults to short-term, single-future thinking, so even good foresight is unusable in the room.

## The guardrails that keep it alive

- ▶ **Statutory routing.** A required input to budget and legislation, not a memo — Finland's and Singapore's lesson made structural.
- ▶ **Centre-of-government seat.** Convening authority no single ministry holds, so foresight can hold the whole state to one shared picture.
- ▶ **The data backbone (Area 31).** An integrated, near-real-time spine so scenarios rest on live, cross-domain signals.
- ▶ **Futures literacy (Area 32).** The human operating system that turns scenarios from reports into reflexes.

### THE HONEST BOTTOM LINE

Every objection on the previous page describes **foresight done badly** — orphaned, captured, ungoverned, unfunded, mistaken for prophecy. The capability described in this report is engineered, point by point, to be none of those things. The risk is real; it is also designable away — and the four foundational areas (1 · 8 · 31 · 32) are precisely the design.

## 1 ÷ N

The decisive KPI: foresight that *changes* a decision over foresight that is merely published

## Traceable

Every agent claim sourced; humans own the judgement

## Structural

Routing, mandate, data and literacy — designed in, not hoped for

MOVEMENT F · THE ENGINE UNDERNEATH

# The Agentic Foresight Architecture

Every area in this report is *run by agents* — but agents need eyes. Foresight is the sensing layer beneath the whole stack: a continuous, agent-run capability that lets a state see disruption while it is still cheap to act on. Twenty-six components, four layers, one engine — built on a dedicated 102-document technology-foresight library.

I Core Agent Fabric

II Data & Knowledge

III Domain Fleets

IV Governance & Wiring


WHY FORESIGHT WORKS · THE FORECASTABLE CURVE

# Disruption is forecastable while it is still cheap to act on

A foresight engine would be a luxury if the future were noise. It is not. The technologies that reshape economies ride *exponential cost-and-capability curves* — and those curves are read-ahead-able. Across 62 technologies, Wright's law (cost falls as a power of cumulative production) forecast progress better than calendar time; the same shape governs solar, batteries and machine-learning compute. The disruption was a surprise only to institutions that were not reading the curve. The agentic architecture exists to read every such curve, continuously, and route what it reads into the budget while the move is still cheap and reversible.

<h2>4–5<sup>x</sup> / yr</h2> <p>growth in frontier-model training compute</p> <p>Epoch AI</p>	<h2>62</h2> <p>technologies where Wright's law beats calendar-time forecasting</p> <p>Nagy &amp; Farmer (MIT)</p>	<h2>81%</h2> <p>of 2023 renewable capacity additions cheaper than the fossil alternative</p> <p>IRENA</p>
--	---	---

REAL COST DECLINE, 2010–2023



Source: IRENA *Renewable Power Generation Costs in 2023*; BloombergNEF.

THE OPTION-VALUE ASYMMETRY

Reading a curve early buys the cheap, reversible move; reacting late buys the expensive, irreversible one. Oxford INET's empirically grounded forecasts found a fast technology-led energy transition is **multi-trillion-dollar net cheaper** than a slow one — the same arithmetic every area in this report describes, now made continuous.

Way, Ives, Mealy & Doyne Farmer (Oxford INET); RethinkX cost-curve disruption framework.

26 · 4

components · layers in the foresight engine

7

reusable agent archetypes do the work

102

primary documents in the evidence library

## THE AGENTIC FORESIGHT ARCHITECTURE AT A GLANCE

## The twenty-six components in one table

#	COMPONENT	LAYER	WHAT ITS AGENTS DO	ANCHOR SOURCE
A1	Agentic Foresight Operating System	I	Orchestrate the fleet; own the shared futures picture	OECD-WEF AI in Strategic Foresight
A2	National Signals Registry & ingestion	I	Curated, deduplicated weak-signal feed	OECD Strategic Intelligence Tools
A3	Scanning agents	I	Continuous horizon scanning & weak-signal detection	SciMON; WISDOM
A4	Forecasting agents	I	Calibrated probabilities on dated questions	Halawi et al; Schoenegger
A5	Scenario-generation agents	I	Recombine signals into stress-tested branches	Ritchey morphological analysis
A6	Simulation & ABM agents	I	Synthetic populations of households & firms	Park et al, Generative Agents
A7	Red-team & adversarial agents	I	Attack every plan and every forecast	CSET; Intl AI Safety Report 2026
A8	Discovery agents	I	Automated tech/science discovery as early signal	Lu et al, The AI Scientist
A9	Briefing agents	I	The two-page weekly anticipatory brief	OECD-WEF; Singapore CSF
A10	Cost-and-capability curve library	II	The forecastable-disruption substrate	Nagy-Farmer; Epoch AI; IRENA
A11	Patent & bibliometric tech-mining	II	Read the frontier before it reaches market	WIPO; OECD Patent Manual
A12	Technology-readiness & roadmapping base	II	Turn a signal into a timed trajectory	NASA TRL; Cambridge IfM
A13	Provenance & the evidence ledger	II	Traceability of every claim the fleet emits	OECD Anticipatory Governance
A14	Human curation & judgement layer	II	Analysts as curators and judges; accountability	Schoenegger, AI-Augmented
A15	AI & frontier-compute fleet	III	Read the curve that bends every other curve	Epoch AI; Stanford HAI Index
A16	Quantum & advanced-computing fleet	III	Price a slow curve with a sharp crypto cliff	WEF; CSET; Lau et al
A17	Engineering-biology & bioeconomy fleet	III	Falling cost curve meets rising risk curve	OECD Bioeconomy; Nat Academies
A18	Energy & clean-tech transition fleet	III	The domain where cost-curve foresight is proven	IEA; IRENA; Oxford INET
A19	Technology-convergence (NBIC) fleet	III	Catch disruption where domains compound	NSF NBIC; OECD STI Outlook 2025
A20	Critical & dual-use / security fleet	III	The geopolitics of the frontier	US CET List; NATO STO
A21	Anticipatory-governance wrapper	IV	Route into rules without freezing innovation	OECD; WEF Agile Governance
A22	Assurance of the agents themselves	IV	Reliability, calibration & red-teaming of the fleet	Intl AI Safety Report 2026
A23	Technology-assessment & parliament	IV	The democratic-legitimacy layer	EP STOA; US OTA legacy
A24	Wiring into budget, STI & procurement	IV	A required input, not a courtesy	OECD STIP; EU RIS3
A25	KPIs & the option-value ledger	IV	Measure lead time & cost avoided, not reports	Singapore CSF; Sitra
A26	The build sequence	IV	First 24 months, talent, the mid-sized state	UNIDO; UK GO-Science

THE SEVEN FORESIGHT AGENT ARCHETYPES

# Seven agents do the work humans cannot scale

The twenty-six components are built from a small, reusable kit of agent archetypes. The evidence that they work is now real, not promissory: language-model forecasters approach the human-crowd aggregate on real questions, an ensemble of models rivals the crowd, and an LLM assistant lifts a human forecaster's accuracy by roughly a quarter. The discipline is equally real — every archetype runs under a named human owner, a calibration score, and an adversarial check.

<p><b>≈ crowd</b></p> <p>LLM forecasters approach the human-aggregate accuracy</p> <p>Halawi et al (Berkeley)</p>	<p><b>+24–28%</b></p> <p>human forecasting-accuracy lift with an LLM assistant (RCT)</p> <p>Schoenegger et al</p>	<p><b>167</b></p> <p>foresight practitioners surveyed on AI-in-foresight</p> <p>OECD-WEF</p>
---	---	--

ARCHETYPE	WHAT IT DOES	EVIDENCE IT WORKS
Scanning	Reads the whole information space; flags weak signals & trend-breaks	SciMON; WISDOM weak-signal detection
Forecasting	Puts calibrated probabilities on dated questions	Halawi; Schoenegger Silicon Crowd
Scenario	Recombines signals into stress-tested narrative branches	Morphological analysis; OECD-WEF
Simulation / ABM	Runs synthetic populations of households, firms, regions	Park et al, Generative Agents (25-agent society)
Red-team	Attacks every plan; hunts the assumption that breaks it	CSET national-power work; Intl AI Safety Report
Discovery	Generates & tests new tech/science as an early signal	Lu et al, The AI Scientist (end-to-end)
Briefing	Compresses the fleet's output into the weekly two-pager	OECD-WEF; Singapore CSF practice

**AGENTS WIDEN THE APERTURE; HUMANS KEEP ACCOUNTABILITY**

No agent output reaches a budget desk or a statute without a named human standing behind it. The curation-and-judgement layer (A14) is not a concession to caution — it is where the architecture's legitimacy lives. Machines flag and model; accountable officials decide and act.

# I The Core Agent Fabric

The engine room (A1–A9): the orchestration layer, the shared signal feed, and the seven agent families that scan, forecast, generate scenarios, simulate, red-team, discover and brief. Build this and a state has a working nervous system, however thin — a continuous loop where a quarterly workshop used to sit.

The fabric's first job is to run foresight *once* for the whole of government rather than thirty times badly. A single operating system owns the shared futures picture; a single signals registry is the evidentiary spine every ministry scans against. On that spine sit the agent families — each cheap to run at the margin, each tireless, none trusted without a human owner. The reframe is that **capacity, not insight, was always the binding constraint**: a handful of analysts could watch only so much. The fabric lifts the ceiling on attention while leaving the ceiling on judgement exactly where it belongs.

#	COMPONENT	THE AGENTIC MOVE & ITS GUARDRAIL
A1	Foresight Operating System	Orchestrates the fleet and owns the one shared futures picture — run once for the whole of government, not thirty times badly. <i>Guardrail</i> : a Chief Foresight Officer accountable for what goes up.
A2	National Signals Registry	A curated, deduplicated weak-signal feed the whole fleet scans against — the shared evidentiary spine, not thirty private spreadsheets.
A3	Scanning agents	Read the full reachable information space daily; flag trend-breaks the moment they cross thresholds. <i>Guardrail</i> : human triage of signal vs noise.
A4	Forecasting agents	Put calibrated probabilities on dated questions; track forecast-vs-outcome. <i>Guardrail</i> : published calibration scores, never a bare number.
A5	Scenario-generation agents	Spin and stress-test narrative branches; pressure-test the canonical national scenarios for staleness.
A6	Simulation & ABM agents	Run policy options across synthetic households, firms and regions to capture behaviour standard models miss.
A7	Red-team agents	Adversarially attack every major plan and forecast; forbidden from authoring the plans they attack.
A8	Discovery agents	Generate and test new technological possibilities — what is becoming feasible before it shows up in patents.
A9	Briefing agents	Compress everything into the PM's two-page weekly anticipatory brief; the human directorate curates what goes up.

## daily

CADENCE · the fabric runs continuously, not as a quarterly set-piece

## 1

PICTURE · one shared view of the futures for the whole of government

## 2 pp

OUTPUT · the weekly anticipatory brief that lands on the desk

### WHY THIS LAYER COMES FIRST

Every later layer plugs into the fabric. Stand up the operating system, the registry and two or three proof-carrying agents, and the data, domain and governance layers become low-cost extensions of a system that already works. Attempt them first and you get isolated tools nobody trusts.

## II The Data & Knowledge Architecture

The living substrate the agents reason over (A10–A14). At its centre is the cost-and-capability curve library — the thing that makes disruption forecastable — flanked by the patent and bibliometric mining layer, the technology-readiness base, the evidence ledger that makes every claim traceable, and the human curation layer where accountability lives.

Watching more signals only changes outcomes if the signals resolve into *timing*. The curve library is what converts a scanning agent's "this is happening" into a forecasting agent's "and it gets cheaper than the incumbent in roughly N years" — the one sentence a budget can act on. Around it, the tech-mining and readiness layers locate where a technology sits on its trajectory, the evidence ledger version-stamps every claim, and the curation layer is where a named human turns the machine's output into an institutional position. This layer is the difference between a fleet that talks and a fleet a finance ministry will believe.

#	COMPONENT	WHAT IT HOLDS & WHY IT MATTERS
A10	Cost-and-capability curve library	Wright's-law experience curves for every tracked technology — the substrate that turns "what's new" into "when it gets cheap." Grounded in Nagy–Farmer, Epoch AI and IRENA cost data.
A11	Patent & bibliometric tech-mining	Reads the frontier before it reaches the market — patent families, preprints, weak-signal topic detection. WIPO, OECD Patent Manual, CWTS Leiden.
A12	Technology-readiness & roadmapping base	Turns a signal into a timed trajectory via TRL and roadmapping. NASA/DoD TRL, Cambridge IfM, UNIDO.
A13	Provenance & the evidence ledger	Every claim the fleet emits is traceable to its source and version — the precondition for routing agent output into power.
A14	Human curation & judgement layer	Analysts as curators and judges. The RCT evidence is that human + LLM beats either alone — so this is a performance layer, not just a control.

**1**  
 shared curve library every domain fleet reads from  
 ENSI design target

**100%**  
 of briefed claims version-stamped to a source  
 ENSI design target

**human+AI**  
 the configuration that beats either alone  
 Schoenegger RCT

**THE SUBSTRATE IS THE MOAT**  
 Agents are commodities; the curated, provenance-stamped, continuously-updated knowledge base they reason over is not. A state that owns its curve library and signals registry owns a compounding asset — every month of operation makes the next forecast sharper.

### III The Domain Agent Fleets

Frontier coverage (A15–A20): specialised fleets that each read one domain's curves, plus a convergence fleet that catches the disruptions no single-domain scan would see. The library behind them is real — 102 documents spanning AI & compute, quantum, biology, energy, convergence and dual-use security.

Each fleet is the Core Fabric pointed at one frontier, fed by that frontier's slice of the curve library. The AI & compute fleet reads the curve that bends every other curve; the energy fleet works the domain where cost-curve foresight is already proven; the biology fleet tracks a falling cost curve against a rising risk curve; the convergence fleet exists precisely because the highest-leverage disruptions fall *between* domains. For a mid-sized industrial state, two of these — AI/compute and energy — carry the most option value and are where a Czech-sized build should start.

#	FLEET	THE FRONTIER SIGNAL IT READS	ANCHOR SOURCE
A15	AI & frontier-compute	Training-compute & capability curves — the curve that bends every other curve	Epoch AI; Stanford HAI
A16	Quantum & advanced computing	A slow capability curve with a sharp cryptographic cliff (harvest-now-decrypt-later)	WEF; CSET; Lau et al
A17	Engineering biology & bioeconomy	A falling cost curve meeting a rising biosecurity risk curve	OECD Bioeconomy; Nat Academies
A18	Energy & clean-tech transition	The most proven cost-curve domain — solar, wind, batteries, electrolyzers	IEA; IRENA; Oxford INET
A19	Technology convergence (NBIC)	Where AI×bio×materials×quantum compound into disruption	NSF NBIC; OECD STI 2025
A20	Critical & dual-use / security	The geopolitics of the frontier — critical-tech lists & economic security	US CET List; NATO STO

#### 6

FLEETS · one per frontier domain, each reading its own curves

#### 1

CONVERGENCE · for the disruptions that fall between domains

#### AI · energy

START HERE · the highest option value for a mid-sized industrial state

#### WHY A CONVERGENCE FLEET EARNS ITS PLACE

The disruptions that broke incumbents in the last decade — AI-designed proteins, software-defined everything, electrified transport — lived between disciplines. A scan organised purely by domain misses them by construction; the convergence fleet (A19) is the deliberate correction.

## IV Governance, Assurance & Wiring

The trust and the route to power (A21–A26). The first three layers build a machine that sees; this layer decides whether anyone may act on what it sees — routing foresight into rules without freezing innovation, assuring the agents themselves, keeping the whole apparatus democratically legible, and hard-wiring it into the decisions that spend money.

A foresight fleet that scans, forecasts and red-teams produces a torrent of anticipatory judgement — and anticipatory judgement is worthless, or worse, unless it is governed on three fronts at once. It must route into rules without freezing the innovation it is reading (A21). The agents producing it must themselves be reliable, calibrated and adversarially tested, because routing an over-confident agent into power is a faster way to make a bad decision (A22). And the whole apparatus must be answerable to a parliament, not just a minister and a vendor (A23). Only then does the wiring (A24) — making the engine a *required* input to the budget and procurement — become safe to switch on.

#	COMPONENT	THE JOB IT DOES
A21	Anticipatory-governance wrapper	Routes foresight into adaptive regulation and sandboxes — acting on the Collingridge dilemma without freezing the future. OECD; WEF Agile Governance.
A22	Assurance of the agents themselves	Reliability, calibration and red-teaming of the fleet — you cannot route an over-confident agent into power. Intl AI Safety Report 2026.
A23	Technology-assessment & parliament	The democratic-legitimacy layer — answerable to a parliament, not just a minister and a vendor. EP STOA; the US OTA legacy.
A24	Wiring into budget, STI & procurement	Makes the engine a <i>required input</i> to the decisions that allocate money — not a courtesy report. OECD STIP; EU RIS3.
A25	KPIs & the option-value ledger	Measures the engine by signal-to-response lead time and cost avoided by acting early — not by reports produced.
A26	The build sequence	The dated 24-month plan, the talent, and the discipline of self-funding each increment from the last.

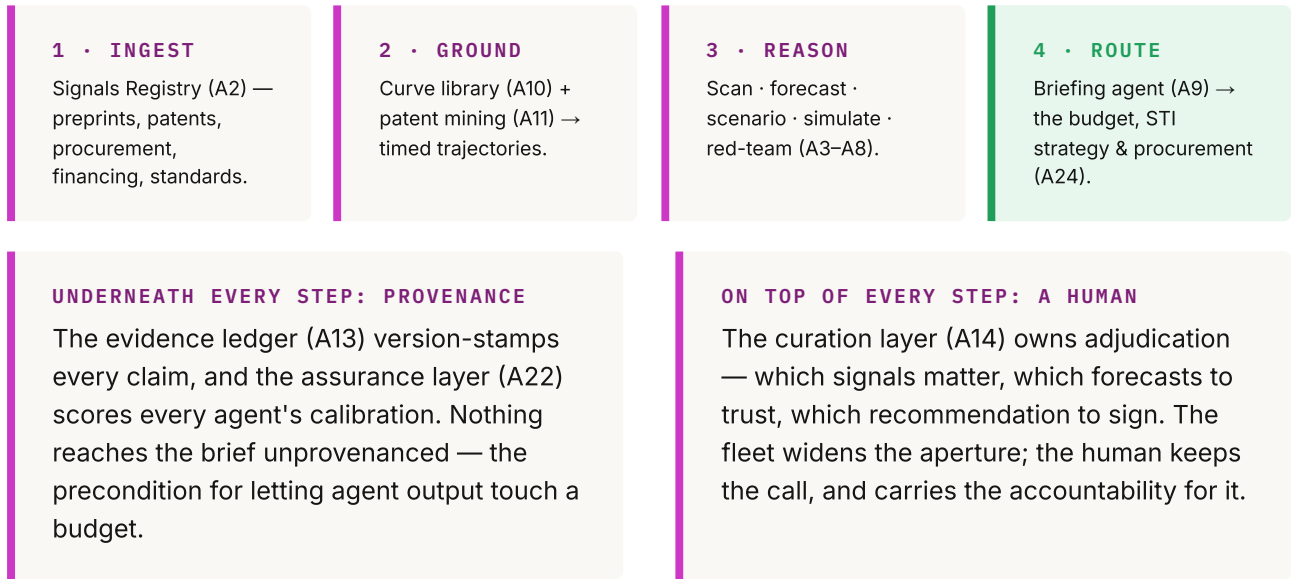
### THE FAILURE MODE THIS LAYER PREVENTS

Anticipatory judgement routed into power *ungoverned* is a faster way to make a bad decision, not a slower one. Assurance and legitimacy are the difference between a clever internal tool and an instrument of the state — which is why the wiring into the budget is switched on last, not first.

THE LIVING SUBSTRATE · HOW THE LOOP RUNS

# Signals in, anticipatory briefs out

The architecture is a single continuous loop. Raw signals are ingested and deduplicated into the registry; the curve library and tech-mining layer turn them into timed trajectories; the agent fleet forecasts, simulates and red-teams; and the briefing layer routes a curated two-pager into the decision — every step provenance-stamped and human-owned. The loop runs daily, where a foresight set-piece used to run once a year.



STAGE	AGENTS INVOLVED	WHAT IT EMITS	THE HUMAN CONTROL
<b>Ingest</b>	Registry & scanning (A2–A3)	Deduplicated weak signals	Triage: signal vs noise
<b>Ground</b>	Curve library, tech-mining (A10–A12)	Timed cost/capability trajectories	Sign-off on the curve fit
<b>Reason</b>	Forecast, scenario, sim, red-team (A4–A8)	Calibrated forecasts & stressed scenarios	Adjudicate disagreement
<b>Route</b>	Briefing & wiring (A9, A24)	The two-page anticipatory brief	Owns & signs the recommendation

## continuous

the loop runs daily, not quarterly

## lead time

the KPI: signal-detection → policy response

## 0


unprovenanced claims permitted in a brief

THE EVIDENCE BASE BEHIND MOVEMENT F

# 102 documents, fifteen angles

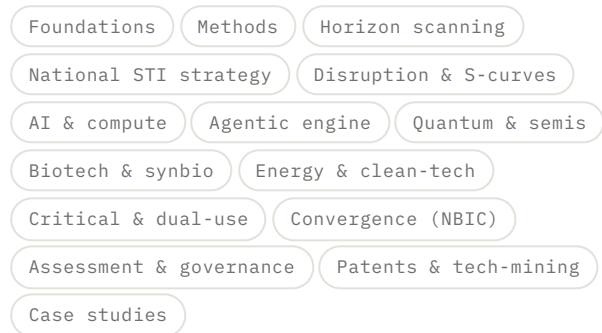
Movement F is built on a dedicated, real, downloaded source library — 102 primary documents across fifteen angles, ~527 MB on disk — pulled from the world's leading foresight and technology institutions. Every claim above is traceable to a named document; none is invented.

FIG. - LIBRARY BY ANGLE CLUSTER



**Fifteen angles, five clusters.** The frontier-domain and agentic-engine streams are what make this a technology-foresight library, not a generic futures one.

THE FIFTEEN ANGLES



REPRESENTATIVE PROVIDERS - INSTITUTIONS

- ▶ **UN** — UNIDO & UNCTAD foresight manuals; WIPO patent-trends.
- ▶ **OECD** — STI Outlook; anticipatory-governance & sandbox toolkits.
- ▶ **EU** — JRC, Commission Strategic Foresight, EP STOA, CORDIS.
- ▶ **Gov** — US OSTP/NSF/OTA, UK GO-Science, Japan NISTEP, Singapore CSF, Sitra, Dubai FF.
- ▶ **IGO** — World Bank, IEA, IRENA, NATO STO, APEC CTF.

REPRESENTATIVE PROVIDERS - FRONTIER & THINK TANKS

- ▶ **Frontier** — Epoch AI, Stanford HAI, Intl AI Safety Report.
- ▶ **arXiv** — Generative Agents, LLM forecasting, the AI Scientist, SciMON.
- ▶ **NBER** — general-purpose-technologies; experience-curve forecasting.
- ▶ **Think tanks** — WEF, RethinkX, CSET Georgetown, Bruegel-class analysis.
- ▶ **Industry** — Deloitte, BloombergNEF cost & trend data.

102

real documents on disk

15 · 5

angles · clusters

~527 MB

downloaded primary-source corpus

THE BUILD SEQUENCE · THE ORDER IS THE STRATEGY

# First 24 months: the foresight engine

A mid-sized EU member state — the running example is the Czech Republic — does not buy foresight with a thousand analysts. It builds the nervous system in the right order, on a modest budget, and proves value against a real decision within the year. The figures below are ENSI design targets, not measured facts: they exist to make the architecture buildable.

WINDOW	WHAT SHIPS	THE PROOF IT MUST PRODUCE
Months 1-6	Foresight OS (A1) + Signals Registry (A2); first scanning & forecasting agents (A3-A4); seed the curve library (A10)	The fleet can sense and forecast at all, with humans curating every output
Months 7-12	Briefing agent (A9) + budget-annex wiring (A24); first option-value ledger entries (A25)	One anticipatory brief lands <i>inside</i> the budget round
Months 13-18	One or two priority domain fleets (for a Czech-sized industrial state: AI/compute A15 + energy A18); assurance & calibration (A22)	A live, dated readiness assessment in two frontier domains
Months 19-24	Red-team & evaluation agents (A7); hardened provenance & evidence ledger (A13); first full option-value portfolio	Year two funded by demonstrated value, not faith

**6-10**  
 senior core team; the agent fleet is the workforce multiplier  
 ENSI design target

**< 12<sub>mo</sub>**  
 to a working minimum-viable architecture  
 ENSI design target

**1**  
 decision-changing brief booked to the ledger in year one  
 ENSI design target

**THE ORDER IS THE STRATEGY**  
 Build the engine and the wiring before the domain fleets; never scale agents faster than the curation team can govern them; and let each increment fund the next from the value the ledger books. For a state with thin fiscal slack, that is not the expensive option — it is the only affordable one.


METHODOLOGY · REPRODUCIBILITY

# How this report was built

This is an analysis-first report built on a real, downloaded source library. The method was deliberate: assemble the evidence base first, derive the 32-area analysis from it, then front-load the report with the argument so a reader who never reaches the catalog still owns the thesis.

- 1 **Assemble the corpus.** We built a library of **188 primary documents across 20 angles** (~928 MB) — every one a real PDF pulled from the original institution or a verified mirror, with its source URL recorded. The angles span what foresight is and how to do it, foresight in and for the state, foresight × AI × agents, the data and statistics literature, and the economics of growth.
- 2 **Derive the analysis.** From that corpus we distilled **32 prioritised areas** of a national foresight capability, ranked by leverage and grouped into **four tiers** — the core, productive capacity, society, and the environment-and-enabler layer. Each area is written as a seven-part operating brief.
- 3 **Front-load the argument.** The report is structured **analysis-first**: a front movement of nine analytical angles carries the full intellectual payload before the catalog. Roughly **40% of the pages are analysis**, placed ahead of the 32-area reference layer — so the understanding forms before the receipts.
- 4 **Wire in the agentic layer.** Every area is specified with the AI agents that run it — the distinctive ENSI move that turns foresight from a quarterly artefact into a continuous capability — grounded in the agentic-foresight and AI angles of the library.

FIGURE · THE CORPUS-TO-REPORT FLOW



**Evidence in, argument out.** The library is the spine; the analysis is the contribution; the catalog is the reference layer that proves it.

**A NOTE ON SOURCING DISCIPLINE**

Facts ENSI owns — 188 sources, 20 angles, 32 areas, four tiers — are stated plainly. Institutional figures and the named foresight bodies (CSF, Sitra, Policy Horizons Canada, Dubai Future Foundation) are attributed to the library. Where a precise statistic was not in hand, the report uses a qualitative claim rather than a fabricated number.


SHOW THE RECEIPTS

# The evidence base: a 188-source library

Every claim in this report stands on a real, downloaded library — **188 primary documents across 20 angles**, each a PDF pulled from the original institution and recorded with its source URL. This is the source mix: where the evidence comes from, and how it is distributed across the questions the report answers.

<h2>188</h2> <p>Primary documents in the library ENSI foresight source library, 2026</p>	<h2>20</h2> <p>Angles, in five thematic clusters Library INDEX</p>	<h2>~928 MB</h2> <p>Of downloaded primary material Library INDEX</p>
--	--	--

FIGURE · SOURCE MIX BY PROVIDER TYPE



- OECD · EU JRC / Commission ~31%
- UN bodies · World Bank · IFIs ~26%
- Academic · arXiv · growth economics ~22%
- National foresight units (CSF, Horizons, Sitra, DFF) ~12%
- Think-tanks (RAND, WEF, Bruegel, CSIS, McKinsey) ~9%

Shares are approximate, derived from the per-angle provider attributions in the library indexes.

WHAT THE MIX BUYS

The library is weighted toward the **institutions that actually run foresight** — OECD, the EU JRC and Commission, the UK Government Office for Science, Policy Horizons Canada, Singapore's CSF, Sitra, the Dubai Future Foundation, the US National Intelligence Council — balanced by the UN and World Bank data literature and the seminal economics of growth.

**NAMED SOURCES IN THE CORPUS**

- OECD
- EU JRC
- UK GO-Science
- Policy Horizons
- Singapore CSF
- Sitra
- Dubai Future Fdn
- US NIC
- UNDP
- UNESCO
- UNECE HLG-MOS
- World Bank
- Eurostat
- PARIS21
- WEF
- RAND
- Bruegel
- CSIS
- McKinsey GI
- Millennium Project
- ČSÚ
- arXiv

**HONEST ABOUT GAPS**

Where a source hard-blocked scripted download (paywalls, bot-gates, Akamai 403s), the library records the block and substitutes an equivalent open-access document — noted per angle.

THE CATALOG · 20 ANGLES

# Every angle, with its document count

The 188 documents are organised into 20 angles across five clusters — what foresight is and how to do it; foresight in and for the state; foresight × AI × agents; the data and statistics literature; and the economics of growth. The counts below are read straight from the library index.

#	ANGLE	CLUSTER	DOCS
01	Foundations of Foresight (What Is Foresight)	Foresight & method	10
02	Foresight Methodologies and Methods	Foresight & method	11
03	Strategic Foresight in Government and Public Policy	For the state	12
04	National and Government Foresight Institutions	For the state	8
05	Horizon Scanning and Emerging Issues	Foresight & method	10
06	Scenario Planning Practice and Frameworks	Foresight & method	9
07	Megatrends and Global Trends Reports	Data & trends	8
08	Foresight and AI (AI-Augmented Futures)	AI & agents	11
09	Agentic Foresight (AI Agents for Futures)	AI & agents	9
10	Anticipatory Governance and Innovation	For the state	10
11	Best Think Tanks — Futures and Foresight Output	For the state	8
12	Data Architecture for Foresight and Decision Intelligence	Data & trends	10
13	Data Governance and What Data to Collect	Data & trends	9
14	National Statistical Offices — Strategy and Modernization	Data & trends	10
15	Official Statistics and Data for Policy	Data & trends	8
16	Theories of Economic Growth	Growth economics	10
17	Drivers of Growth — What Makes Countries Grow	Growth economics	8
18	Economic Complexity and Industrial Policy	Growth economics	10
19	Technology and STI Foresight	AI & agents	9
20	Foresight for Resilience, Risk and Long-Term Decisions	For the state	8

**188**

Documents, summed across  
20 angles

**5**

Thematic clusters

**12**

Largest angle: Strategic  
Foresight in Government

**8**

Smallest angles, each fully  
populated

HOW TO READ THE CATALOG

Each angle folder carries its own index listing every document's *provider, type, source URL, file and status* — plus a *Missing* note recording any source that hard-blocked download and the open-access substitute used. The library is built to be audited, not just cited.

## GLOSSARY · THE WORKING VOCABULARY

## The foresight lexicon, in plain terms

Seven terms recur throughout this report. Each names a discipline, not a buzzword — a specific way of using the future to make a better decision in the present.

### Horizon scanning

The continuous, structured search for early signals of change — weak signals, trend-breaks, emerging issues — before they are obvious. The radar of a foresight system.

### Weak signals

Faint, early indicators of a change that is not yet established — the tremor before the shift. The discipline is telling signal from noise by tracking which ones strengthen.

### Scenario planning

Building a small set of distinct, plausible futures to stress-test strategy against — not to predict which arrives, but to find the choices that hold up across all of them.

### Anticipatory governance

Governing so that anticipation is built into how decisions are made — routing foresight into the budget, legislation and strategy cycles, rather than reacting after the fact.

### Backcasting

Fixing a desired future and working backwards to the dated sequence of decisions required to reach it — turning a long horizon into present action.

### Futures literacy

The human capability — recognised as a competence by UNESCO — to use the future to see the present differently. The operating system that turns scenarios from reports into reflexes.

### Three Horizons

A frame separating today's maturing system (H1), the emerging future already present in pockets (H3), and the contested transition between them (H2) — where most strategy mistakes are made.

### ONE ENGINE, MANY METHODS

These are not rival techniques but layers of one stack: scan for signals, build scenarios from them, backcast to decisions, and keep people literate enough to act. The engine is the same; only the domain changes.

# Stop being surprised by a future you could have seen.

A country does not need better forecasts. It needs a standing foresight capability — scan, model, stress-test, route to decision — run continuously by AI agents, because the cost of acting early is a fraction of the cost of reacting late. Compounded across 32 domains and across decades, that gap is the difference between a country carried by the future and one that authors it.