The Foresight Handbook

Centre for Administrative Service

Table of Contents

Right-click and select 'Update Field' to generate the Table of Contents.

Right-click the table above and select “Update Field” to generate page numbers.

Preface

This Handbook provides guidance on the practice of foresight in government planning. It is a systematic treatment of the functions, processes, and craft of foresight: the disciplines by which a government explores possible futures to inform present decisions. The Handbook covers horizon scanning, scenario methods, trend analysis, expert elicitation, and the institutional arrangements that make foresight effective. It is the companion to the Planning Manual and assumes the reader has mastered the doctrine of that volume or its equivalent.

The Handbook is written for the foresight practitioner as a dedicated professional category, and for those who discharge foresight functions under other titles. The techniques and principles it covers are drawn from comparative foresight practice across many jurisdictions and traditions; they are not the property of any single polity.

The Handbook is organized in five parts. Part I establishes the foundations of foresight statecraft. Part II treats horizon scanning in operational detail. Part III treats scenario methods. Part IV treats foresight methods including expert elicitation, cross-impact analysis, trend analysis, and visioning. Part V addresses foresight in practice: designing and conducting exercises, producing foresight products, and institutionalising foresight.

Corrections and proposed revisions should be submitted through the standard review process.

Chapter 1 — What Foresight Is

Part I: Foundations of Foresight — The discipline of exploring possible futures to inform present decisions, and the foresight practitioner's place within it.

1.1 Definition

Foresight is the structured exploration of possible futures to inform present decisions. The definition carries four load-bearing elements: structure, exploration, possibility, and present action. Structure distinguishes foresight from speculation, requiring defined methods, documented assumptions, and reviewable outputs. Exploration, possibility, and present action complete the definition: the discipline opens the range rather than closing it, preserves that range as a set rather than a single projection, and is judged by its effect on choices made today rather than on the accuracy of any one scenario. Each element excludes a corresponding simplification: speculation in place of analysis, projection in place of exploration, a single forecast in place of the set, and the deferral of present choice to an indefinite future.

Foresight sits within government as an analytical function distinct from policy analysis, planning, and intelligence. Policy analysis evaluates the options available to address a defined problem. Planning translates chosen objectives into commitments over a defined horizon. Intelligence assesses the current and near-term operating environment. Foresight addresses the longer-term conditions under which policy, planning, and intelligence must operate, producing structured assessments of the future and leaving the choice of objectives and means to the political and strategic authorities.

Foresight does not predict; it prepares. The discipline produces no claim about which future will materialise. It produces a map of the territory the planner must navigate, with the contours of plausibility marked and the conditions that would move the planner from one part of the map to another specified. The product is preparation: a plan that holds under each of the futures considered, a monitoring system that detects which future is emerging, and a revision process that adapts when conditions shift. A planner who asks whether foresight got the future right has misunderstood the product, which is judged by the quality of preparation it enables.

The range of possibilities is foresight's distinctive output. The discipline preserves the range as a set of distinct, internally consistent futures, each plausible on its own terms rather than as a variant of a central case. The range is bounded by what the evidence supports, what the methods produce, and what the planning question requires. A planner who commissions foresight receives, in return, a structured account of the conditions the plan must withstand. The account is the discipline's contribution to the plan; the choices the planner makes on its basis remain the planner's responsibility.

1.2 The Distinction from Forecasting

Forecasting and foresight address uncertainty with different purposes, methods, and outputs. Forecasting projects what is likely to happen, taking a measured trend or a calibrated model and extending it into the future with a probability distribution where the method supports one. Foresight explores what could happen, constructing multiple futures, each plausible on its own terms, and assessing the implications of each. The forecaster's product is a projection; the foresight practitioner's product is a set of scenarios. The two products serve different decisions, and substituting one for the other at the wrong point in the planning cycle compromises both.

The difference in purpose follows from the difference in object. Forecasting serves decisions that depend on the most likely outcome: revenue projections for the next fiscal year, demand for a regulated service, the cost of a defined commitment. Foresight serves decisions that must hold across a range of outcomes: infrastructure investments with thirty-year lives, sectoral strategies exposed to technological disruption, institutional reforms whose effects compound over decades. The forecaster reduces uncertainty to a point estimate so that the decision can be optimised; the foresight practitioner preserves uncertainty as a range so that the decision can be made robust. The two purposes sit at different points in the planning process, and a mature planning authority maintains the capability for both.

The difference in method follows from the difference in object. Forecasting relies on statistical extrapolation, econometric modelling, and structured expert judgement where data are sufficient and the underlying processes are stable. Foresight relies on horizon scanning, scenario construction, driver mapping, and expert elicitation where data are partial and the underlying processes are in flux. The two methods draw on different evidence and different forms of expertise. Forecasting fails when the future departs from the model; foresight fails when the constructed range excludes the future that materialises.

Planning requires both. A plan that rests on forecasting alone treats the baseline as a prediction and leaves the plan exposed to any departure from it. A plan that rests on foresight alone treats the range as a substitute for commitments and leaves the plan without a defensible central case. Mature practice sets the plan against a forecast baseline, then tests the plan's resilience against a foresight range, with the two disciplines sequential and complementary. Where the planning authority maintains both capabilities in-house, the disciplines are kept analytically distinct even where they draw on common data and expertise.

1.3 Foresight as a Discipline

Foresight is a discipline in the strict sense: a body of methods, a set of institutional arrangements, and a recognised professional category. The methods are documented, taught, and refined through practice. The institutional arrangements are observable across mature jurisdictions, from dedicated units at the centre of government to parliamentary committees with foresight mandates. The professional category has its own journals, associations, degree programmes, and career paths. A planner who treats foresight as an occasional exercise or as a generic form of thinking about the future is not practising the discipline, and the products of such an exercise will not withstand scrutiny.

The methods of foresight include horizon scanning, scenario construction, driver mapping, expert elicitation, and cross-impact analysis. Horizon scanning collects and assesses signals of change across defined source types on a defined schedule. Scenario construction builds internally consistent narratives of plausible futures along the axes of critical uncertainty. Driver mapping identifies the forces that cause observed trends and the relationships among them, while expert elicitation structures the judgement of specialists whose knowledge cannot be reduced to data. Cross-impact analysis traces the cascade of effects through the driver network, identifying leverage points and feedback loops.

The institutional arrangements vary, but share common features. Singapore's Centre for Strategic Futures, located within the Prime Minister's Office strategy apparatus, conducts horizon scanning and scenario development as a continuous function. The United Kingdom's Government Office for Science, under the Chief Scientific Adviser, produces thematic foresight reports that feed cross-government strategy. Finland's Committee for the Future provides a parliamentary forum for foresight, linking the discipline to legislative scrutiny. France's Commissariat Général du Plan, during its operational period, maintained foresight as a standing input to the indicative planning process, and the European Commission's Joint Research Centre supports foresight across the Union's policy domains.

The professional category has consolidated over the past three decades. Foresight practitioners work in dedicated units, in strategy functions, and as external consultants to government. The discipline has its own degree programmes, professional associations, and methodological literature, and its practitioners are trained in the methods named above and in the facilitation skills that scenario workshops require. Practitioners are also trained in the discipline's limits: when to recommend forecasting rather than scenarios, when to defer to sectoral experts, and when to acknowledge that the evidence will not support a definitive range. the doctrine treats foresight as a discipline in this full sense: a specialism with standards the planner must meet or commission.

1.4 The Horizon Question

Foresight is conducted over horizons of ten, twenty, and thirty years, and occasionally longer. The horizon distinguishes foresight from the planning cycle: the budget horizon is one year; the plan horizon is typically five; foresight extends two to six times the plan horizon. The longer horizon reflects the time over which the forces shaping the plan — demographic change, technological transition, climate shift, geopolitical realignment — unfold and produce their effects. A planner who limits foresight to the plan horizon tests the plan against the conditions of its own assumptions. The forces named above operate on timescales that exceed any single plan period, and their effects accumulate across successive plans.

The horizon is set by the question. A planner examining workforce supply for a regulated profession sets the horizon at the time required to train a new entrant, ten to fifteen years from school to senior practice. A planner examining the implications of climate transition for coastal infrastructure sets the horizon at the design life of the assets, thirty to fifty years. A planner examining the fiscal implications of demographic ageing sets the horizon at the period over which

the demographic shift will unfold, forty years or more. The horizon is the time over which the decision's consequences run, and the planner sets it by examining the decision.

The long horizon matters because the decisions the plan shapes have long lives. Infrastructure assets depreciate over thirty to fifty years. Sectoral strategies compound over a generation. Constitutional and institutional reforms operate indefinitely, and a plan that looks only five years ahead commits the jurisdiction to decisions whose consequences extend far beyond the window examined. Foresight at ten, twenty, or thirty years brings those consequences into view while there is still time to choose.

The trade-off between horizon and certainty is the central constraint. The further ahead the planner looks, the wider the range of plausible futures and the less any single future can be defended. A ten-year horizon may support a central case surrounded by variants; a thirty-year horizon may support only a set of qualitatively distinct futures with no defensible central case. The planner does not respond to the trade-off by shortening the horizon; the planner responds by changing the method. As the horizon lengthens, the discipline shifts from forecasting toward scenario construction, from point estimates toward ranges, and from optimisation toward robustness.

1.5 Foresight and the Plan

Foresight informs the plan through three mechanisms: the assumptions, the risks, and the revision process. The mechanisms are sequential in the planning cycle and continuous across it. The plan's assumptions state the conditions the plan takes as given over the horizon. The plan's risks identify the conditions under which the assumptions may prove wrong, and the revision process adapts the plan when conditions change. The three mechanisms are interlocking, and each requires a counterpart in the planning institution: an assumptions log, a risk register, a monitoring system, and a scheduled revision cycle.

The first mechanism is the assumptions. Foresight produces a range of plausible conditions — demographic trajectories, growth paths, technology baselines, geopolitical environments — and the plan's assumptions are selected from that range. Selection is a judgement rather than a prediction: the planner chooses assumptions that are plausible, internally consistent, and sufficient to support the plan's commitments. The assumptions are documented, dated, and traceable to the foresight products on which they rest, so that a plan whose assumptions cannot be traced to a foresight exercise is a plan whose foundations have not been examined. Traceability allows the assumptions to be revisited when the conditions on which they rest are revised, and it permits auditors and reviewers to test the rigour of the selection.

The second mechanism is risk treatment. The scenarios produced by foresight identify the conditions under which the assumptions may fail, and these conditions are treated as strategic risks. Each risk is assessed for likelihood and impact, and each is addressed through mitigation, transfer, acceptance, or avoidance. A scenario in which a key technology matures faster than assumed becomes a risk to a plan that commits to a specific technology path, addressed through phased commitments, parallel investment, or contractual flexibility. Foresight brings the risk into the plan's treatment so that the plan's exposure is explicit and managed.

The third mechanism is plan revision. Foresight produces indicators that distinguish one scenario from another: the early signals that mark which future is emerging. These indicators are added to the monitoring system, and the monitoring results feed the revision process. When a scenario begins to materialise, the plan's assumptions are revisited and, where necessary, revised. The three mechanisms together constitute the planner's response to uncertainty: foresight provides the range, the plan provides the commitment, and the revision process provides the adaptation.

1.6 Chapter Summary

Foresight is the structured exploration of possible futures to inform present decisions. The discipline produces a range of plausible futures rather than a single projection, and is judged by its effect on present choices rather than by the accuracy of any one scenario. Foresight does not predict; it prepares.

Foresight is distinguished from forecasting by its purpose, method, and output. Forecasting projects the most likely future for decisions that can be optimised against it; foresight constructs a range of futures for decisions that must hold across them. Planning requires both: the forecast baseline to set the central case, and the foresight range to test it. The discipline rests on documented methods — horizon scanning, scenario construction, driver mapping, expert elicitation, cross-impact analysis — and on institutional arrangements observable across mature jurisdictions, from Singapore's Centre for Strategic Futures to the United Kingdom's Government Office for Science to Finland's Committee for the Future.

The foresight horizon extends well beyond the plan horizon, typically to ten, twenty, or thirty years, and is set by the question rather than by convention. As the horizon lengthens, the method shifts from forecasting toward scenario construction, and from optimisation toward robustness. Foresight informs the plan through three mechanisms: the assumptions, selected from the foresight range; the risks, identified from the scenarios; and the revision process, triggered by the indicators foresight specifies. The three together constitute the planner's response to uncertainty.

The relationship between foresight and the plan is maintained across the planning cycle, not only at the plan's start. Foresight at the start of the cycle establishes the assumptions and the risks; foresight during the cycle updates the indicators and the scenarios; foresight at the revision point tests whether the emerging future remains within the constructed range. The discipline is a continuous function, and the planning authority that treats it as a one-off input has not yet built the relationship the handbook describes.

Chapter 2 — The Foresight Function in Government

The foresight function is organised through institutions.

2.1 The Function

The foresight function is the institutional capacity by which a government systematically explores possible futures to inform present decisions. It discharges this capacity through four activities: it produces foresight products, it maintains the scanning capability, it advises the strategy function, and it contributes to the planning cycle. Each activity is distinct and each requires dedicated resources. A function that produces occasional reports but does not scan has a partial capability; a function that scans but does not advise has a library rather than a function. The four together constitute the foresight function as a governmental capability.

Production of foresight products is the most visible activity. The products include horizon scan briefings, emerging issues reports, scenario sets, driver maps, and strategic foresight studies on defined themes. Each product has a defined audience, a defined format, and a defined place in the decision process. Production is conducted on a schedule, to a standard, and for a stated purpose, rather than as a response to ad hoc requests. A function whose products are produced only when commissioned has ceded its work programme to its customers and cannot maintain the systematic character foresight requires.

Maintaining the scanning capability is the continuous activity that underpins production. Horizon scanning collects, filters, assesses, and prioritises signals of change on a standing basis, drawing on the scientific literature, the technical press, government reports from other jurisdictions, international organisation assessments, and expert elicitation. The capability requires dedicated staff, defined sources, a maintained scan log, and a review cycle. Scanning is the most resource-intensive element of the function and the most easily neglected, because its outputs accumulate gradually and have no single dramatic moment of impact. A scanning capability allowed to lapse cannot be rebuilt quickly when a crisis demands it.

Advising the strategy function and contributing to the planning cycle are the activities through which foresight becomes operational. The strategy function sets the direction of the plan; foresight supplies the conditions the strategy must address. The planning cycle produces and revises the plan; foresight supplies the assumptions, risks, and early warning the cycle requires. Without these two connections, the foresight function produces scholarship rather than governmental capability. The relationships are maintained through structured handoffs, treated in section 2.3 and elaborated throughout the Handbook.

2.2 Institutional Forms

The foresight function is organised in four principal forms across mature jurisdictions: the central foresight unit, the distributed model, the parliamentary committee, and the advisory body. The forms vary in their location within the architecture of government and in their relationship to

political authority. No single form is correct; each is appropriate to the constitutional arrangements, the administrative tradition, and the scale of the jurisdiction in which it operates. The choice among them is governed by these conditions, and the comparative examples in section 2.5 illustrate how each works in practice.

The central foresight unit is a dedicated body within the executive, typically located at the centre of government. Singapore's Centre for Strategic Futures, within the Prime Minister's Office Strategy Group, is the leading example. The advantage of the central unit is proximity to decision-making: foresight products reach the highest level of the executive directly, and the unit can commission work across government. The risks are capture and fragility. A central unit may serve the immediate preoccupations of the government of the day rather than the longer horizons foresight requires, and a change of government or a reorganisation can dissolve the capability.

The distributed model locates foresight capacity across sectoral ministries, with coordination by a central body. The advantage is depth: sectoral foresight analysts have the domain knowledge a central unit cannot maintain, and the capacity is harder to dissolve because it is embedded across the administration. The disadvantage is fragmentation. Without strong coordination, sectoral foresight produces inconsistent methods, duplicated effort, and gaps at the intersections between sectors. The distributed model requires a coordinating body, shared standards of method, a common scan log, and a community of practice that meets regularly to maintain consistency.

The parliamentary committee form provides a legislative forum for foresight. Finland's Committee for the Future, established in 1993, is the principal example: the committee receives the Government's Report on the Future, conducts its own assessments, and provides a parliamentary input to the foresight function. Its strength is legitimacy, as foresight conducted through a parliamentary committee carries the authority of the legislature and connects directly to the political process; its limit is distance from the executive's planning machinery, as a parliamentary committee can advise and assess but does not direct the executive's plan. The advisory body form, exemplified by France Stratégie, provides expert foresight to the executive from outside the line administration. The value here is independence; the constraint is the need for a formal channel into the planning process, which must be maintained by both sides.

2.3 The Foresight Cycle

The foresight cycle is the rhythm by which the foresight function produces its work and connects to the planning cycle. The planning cycle, treated in the Planning Manual, comprises diagnosis, strategy, production, approval, execution, monitoring, and revision. The foresight cycle runs alongside this cycle and feeds into it at defined points. Alignment is deliberate. Foresight that is not synchronised with the planning cycle produces work that arrives too late to shape the plan or too early to be remembered when the plan is drafted.

The foresight cycle has four phases. Scoping defines the question, the horizon, the scope, and the methods; collection gathers signals, data, and expert input; analysis identifies the key drivers, constructs scenarios, and assesses implications; and dissemination communicates the foresight product to the strategy function and the planning authority. The four phases are treated at the general level in Chapter 6 of the Planning Manual; the operational detail for each is treated

throughout. The cycle is iterative, and the iteration is itself a discipline: each pass refines the question, the evidence, and the product.

The alignment with the planning cycle requires that foresight exercises be timed to feed the plan's diagnosis and strategy phases. Foresight conducted during diagnosis supplies the conditions the diagnosis must consider. Foresight conducted during strategy supplies the scenarios against which the strategy is tested. Foresight conducted during production is too late to shape the plan and can only comment on what has been drafted. The foresight function's work programme is therefore set by the planning calendar, not by the function's own preferences, and the function maintains the calendar explicitly to govern its programme.

Continuous scanning runs independently of the cycle. The scan is a standing capability, conducted on a fixed schedule of source coverage and review, not a phased exercise. It feeds into each foresight exercise as it is conducted and produces its own standing outputs — the scan briefing and the emerging issues report — on a regular schedule regardless of where the planning cycle stands. The relationship between the continuous scan and the phased foresight exercise is one of the institutional design questions treated in Chapter 5. The scan supplies the raw material; the exercises produce the synthesised products; the cycle ensures both connect to the plan.

2.4 The Foresight Unit's Composition

The foresight unit comprises four roles: the foresight analyst, the domain expert, the facilitator, and the communicator. In a small unit, one person may discharge several roles; in a large unit, the roles are distinct. The composition reflects the range of competencies the function requires. A unit staffed only by generalists lacks substantive depth; a unit staffed only by domain experts lacks method; a unit without facilitation capacity cannot conduct the workshops on which much foresight work depends; a unit without communication capacity cannot disseminate its products to the audiences that must use them.

The foresight analyst is the methodological core of the unit. The analyst conducts horizon scanning, constructs scenarios, designs expert elicitation, and produces the analytical content of foresight products. The analyst's craft is treated in Chapter 4. The domain expert supplies the substantive knowledge the analyst's methods require: the energy expert for an energy transition foresight exercise, the demographic expert for a population study, the security expert for a geopolitical assessment. Domain experts may be permanent members of the unit, seconded from sectoral ministries for defined periods, or engaged as external consultants for specific exercises.

The facilitator designs and runs the workshops through which much foresight work is conducted. Scenario workshops, expert elicitation panels, and driver mapping sessions all require skilled facilitation. The facilitator's craft is distinct from the analyst's: the facilitator manages the process, ensures all voices are heard, and produces the record of the session, while the analyst supplies the content and the synthesis. The communicator translates foresight products into the formats the strategy function and the planning authority can use, and the craft is the judgment required to convey uncertainty without undermining the product's authority and to convey implications without overstepping the analytical record.

The unit's relationship to the statistics office and the planning authority is structural. The statistics office is the authoritative source for quantitative data; the foresight unit draws on it and does not duplicate its work. The planning authority is the principal customer for foresight products; the unit's work programme is set in consultation with the planning authority and governed by the planning calendar. The relationship to the intelligence function, which also conducts scanning and assessment, is treated in Chapter 3. Maintaining clear relationships with these adjacent functions prevents duplication, clarifies responsibility, and ensures the foresight unit's products reach their intended users.

2.5 Comparative Examples

Singapore's Centre for Strategic Futures, established in 2009 and operating within the Prime Minister's Office Strategy Group, conducts horizon scanning, scenario development, and strategic foresight studies as a continuous function. The CSF's work is characterised by integration with the centre of government and by emphasis on whole-of-government scenarios. The Strategy Group provides the channel through which foresight products reach the highest level of the executive, and the CSF contributes to foresight capacity across the Singapore public service through training and the publication of methodological guidance. The CSF's contribution to the practice is the demonstration that a central foresight unit, properly resourced and properly located, can sustain a continuous foresight capability that feeds directly into national strategy.

The United Kingdom's Government Office for Science, operating within the Department for Science, Innovation and Technology, produces foresight reports on defined themes through its Foresight programme. The Government Chief Scientific Adviser, who heads the office, provides a channel to the Prime Minister and the Cabinet. The UK model is characterised by thematic depth: each Foresight project assembles a large expert network and produces a substantial evidence base over a period of two to three years. The advantage of the model is rigour; the challenge is timeliness, as the long project cycle can leave the function slow to respond to fast-moving developments. The contribution to the practice is the demonstration that foresight can be conducted as a series of substantial evidence-based projects with the scientific advisory system as its channel.

Finland's Committee for the Future, established in 1993, is a standing committee of the Parliament that receives the Government's Report on the Future, conducts its own assessments, and provides a parliamentary input to the foresight function. Finland also operates a distributed foresight function across ministries, coordinated through the Prime Minister's Office, and the combination of parliamentary and executive foresight is distinctive. The Finnish model's contribution to the practice is the demonstration that foresight can be conducted with the legislature as a partner, not only as an audience, and that this arrangement strengthens the legitimacy of long-horizon thinking in the political process. France Stratégie, the Commissariat général à la stratégie et à la prospective, advises the Prime Minister on strategic and foresight questions from outside the line administration; its contribution is the demonstration that an independent advisory body, properly mandated and properly connected, can supply foresight to the executive with intellectual depth that an in-house unit may find difficult to sustain.

The European Commission's Joint Research Centre, through its foresight team, supports the Commission's policy work with foresight methods across the Union's policy domains. The JRC's contribution to the practice includes the development of methods for participatory foresight and the conduct of foresight exercises that span multiple member states. The OECD's strategic foresight work, conducted through the OECD Secretariat, consolidates practice across member states and produces comparative guidance. Together with the bodies above, these institutions constitute a community of practice from which the doctrine in the doctrine is drawn. Each contributes particular strengths: the CSF its integration with the centre; the GO-Science its thematic rigour; the Finnish committee its legislative engagement; France Stratégie its independence; the JRC its methodological development; the OECD its comparative consolidation.

2.6 Chapter Summary

The foresight function is the institutional capacity by which a government systematically explores possible futures. It discharges this capacity through four activities: producing foresight products, maintaining the scanning capability, advising the strategy function, and contributing to the planning cycle. A function that omits any of the four is incomplete. Production without scanning is unsustained; scanning without advisory is unconnected; advisory without contribution to the cycle is uninformed; contribution without production has nothing to contribute.

The function is organised in four principal forms: the central foresight unit, the distributed model, the parliamentary committee, and the advisory body. Each has strengths and weaknesses, and the choice among them is governed by the constitutional arrangements, the administrative tradition, and the scale of the jurisdiction. The central unit offers proximity to decision-making but risks capture and fragility; the distributed model offers depth but requires strong coordination; the parliamentary committee offers legitimacy but is distanced from executive planning; the advisory body offers independence but requires a maintained channel into the planning process.

The foresight cycle runs alongside the planning cycle and feeds into it at defined points. Foresight conducted during the plan's diagnosis and strategy phases shapes the plan; foresight conducted during production is too late. The cycle's alignment with the planning calendar is the institutional mechanism by which foresight becomes operational. Continuous scanning runs independently of the cycle and supplies the standing outputs that hold the function together between major exercises.

The foresight unit comprises the analyst, the domain expert, the facilitator, and the communicator. The composition reflects the range of competencies the function requires. The unit's relationships to the statistics office, the planning authority, and the intelligence function are structural and must be maintained to prevent duplication and ensure the unit's products reach their users. The comparative examples — Singapore's CSF, the UK's GO-Science, Finland's Committee for the Future, France Stratégie, the European Commission's JRC, and the OECD's strategic foresight work — illustrate the range of practice from which the doctrine is drawn.

Chapter 3 — Foresight and Adjacent Disciplines

Foresight is related to but distinct from four adjacent disciplines.

3.1 Foresight and Forecasting

Foresight and forecasting share data sources and parts of their analytical apparatus but differ in purpose, method, and output. Forecasting is the projection of likely future conditions from observed data, producing a point estimate with a stated confidence interval. Foresight is the structured exploration of plausible futures, producing a set of scenarios with implications for present decisions. The two disciplines are not in competition; each addresses a question the other cannot. The mature planning authority maintains capacity in both, and the planner draws on each for the purpose it serves.

The distinction in purpose determines when each discipline is used. Forecasting serves questions that admit a probabilistic answer: what will the population be in five years, what will revenue be next fiscal year, what will demand for a service be at the end of the plan period. Foresight serves questions that do not: how might the geopolitical environment evolve over twenty years, what are the implications of a technological discontinuity, what conditions would render the plan's assumptions untenable. Forecasting produces the baseline against which the plan is constructed; foresight produces the conditions against which the baseline is tested. A planner who substitutes forecasting for foresight produces a plan that assumes a single future; a planner who substitutes foresight for forecasting produces a plan without a defensible central case.

Complementarity operates through the plan's structure. The baseline, treated in Chapter 9 of the Planning Manual, is constructed from forecasts produced by the statistics office and the economic forecasting authority. The assumptions on which the baseline rests are tested by foresight: the conditions that would invalidate the forecasts, the trends that could break, the shocks that could disturb the trajectory. The risk register, informed by foresight, identifies the conditions the plan must be prepared to address. Where the two disciplines are conducted by the same body, the discipline of distinguishing their outputs must be maintained; where they are separated, the handoff must be deliberate.

Several jurisdictions illustrate the institutional arrangements. The United Kingdom's Office for Budget Responsibility produces the fiscal and economic forecasts that inform the budget, while the Government Office for Science produces the foresight reports that inform longer-horizon strategy. Singapore's Ministry of Trade and Industry produces economic forecasts, while the Centre for Strategic Futures conducts foresight on the longer horizon. The separation is not universal: in some jurisdictions the same body produces both. What matters is that the discipline of distinguishing the outputs is maintained, and that the planner uses each for the purpose it serves.

3.2 Foresight and Intelligence Analysis

Intelligence analysis is the assessment of current and near-term threats to the state's security and foreign policy interests; foresight is the exploration of medium- and long-term conditions that will shape the planning environment. The two disciplines share methods — horizon scanning, expert elicitation, structured assessment — but operate on different horizons and serve different decisions. Intelligence analysis informs immediate and near-term decisions about security, foreign policy, and protective action. Foresight informs medium- and long-term decisions about planning, strategy, and resource allocation. The disciplines are complementary where their horizons overlap and distinct where they do not.

The open-source intelligence connection is the principal bridge between the two disciplines. Open-source intelligence is the systematic collection and assessment of publicly available information, and horizon scanning in foresight uses substantially the same sources and similar methods. The intelligence function collects open-source material to assess current threats; the foresight function collects it to detect emerging change. The same signal — a technological development, a political shift, a demographic movement — may be relevant to both functions, but its assessment and its use differ. The intelligence function assesses the signal's implications for current security; the foresight function assesses the signal's implications for the planning horizon.

Shared methods create opportunities for collaboration that the mature jurisdiction exploits. The intelligence function's scanning capability may detect signals relevant to foresight; the foresight function's scanning may detect signals relevant to intelligence assessment. Exchange of sources, methods, and assessments between the functions can improve the quality of both. The collaboration is most productive where the functions remain institutionally distinct, each maintaining its own standards, its own purpose, and its own customer. A foresight function absorbed into the intelligence apparatus loses its long-horizon character; an intelligence function that absorbs foresight dilutes its focus on current threats.

The comparative arrangements confirm the distinction. The United States separates the intelligence community from the foresight apparatus, with the National Intelligence Council's Global Trends reports sitting at the boundary and the intelligence agencies producing current assessments. The United Kingdom separates the Joint Intelligence Committee from the Government Office for Science's foresight work. Singapore's Centre for Strategic Futures operates alongside the security apparatus within the Prime Minister's Office strategy group. The arrangement is governed by the jurisdiction's administrative tradition and the scale of its intelligence capability, but in each case the two disciplines serve different decisions and produce different products.

3.3 Foresight and Policy Analysis

Policy analysis is the assessment of options for current decisions; foresight is the exploration of the futures those options must address. The two disciplines are sequential in the planning process: foresight identifies the conditions the policy must address, and policy analysis assesses the options for responding to them. A policy analysis conducted without foresight addresses the present as if it were permanent; a foresight exercise conducted without policy analysis produces futures

without action. Planning requires both, and the relationship between them is one of the principal channels through which foresight becomes operational.

The sequential relationship operates through the planning cycle. Foresight identifies the conditions that will shape the policy environment over the horizon: the demographic shifts, the technological developments, the economic trends, the geopolitical changes. Policy analysis assesses the options for responding to those conditions: the regulatory choices, the investment choices, the institutional choices. The sequence is not strictly linear; foresight may be revised in light of policy analysis, and policy analysis may identify questions that require further foresight. The direction of flow, however, is from foresight to policy analysis: the futures first, the options second.

The distinction matters for the organisation of the work. A policy analysis unit that conducts its own foresight risks producing futures that confirm its preferences; a foresight unit that conducts its own policy analysis risks producing advocacy rather than exploration. The mature practice separates the functions while ensuring their connection. The foresight function produces the futures; the policy analysis function assesses the options; the planning authority integrates the two into the plan. The handoff between foresight and policy analysis is one of the critical interfaces in the planning process, and its quality determines whether foresight informs the plan or sits beside it.

The comparative experience confirms the institutional separation. France Stratégie advises the Prime Minister on strategic and foresight questions, while policy analysis is conducted within the sectoral ministries and through the Conseil d'État's advisory functions. The United Kingdom's Government Office for Science produces foresight, while policy analysis is conducted by departments and by the Cabinet Office strategy unit. Singapore's Centre for Strategic Futures produces foresight, while policy analysis is conducted by the line ministries. The arrangement is consistent across these jurisdictions: foresight and policy analysis are organisationally distinct, each maintaining its standards and its purpose, and connected through the planning cycle.

3.4 Foresight and Risk Assessment

Risk assessment is the identification and assessment of risks to the plan; foresight is the identification of the conditions that could create risks. The two disciplines are related but distinct. Risk assessment works from known risks — the events and conditions that have been identified as threats — and assesses their likelihood and impact. Foresight works from the range of possible futures and identifies the conditions that could produce risks not yet visible. Risk assessment is concerned with the known; foresight is concerned with the conditions that could expand or transform the known.

The foresight input to strategic risk assessment is one of the principal channels by which foresight becomes operational. The risk register, maintained by the planning authority, lists the risks the plan must be prepared to address. A risk register informed only by known risks is a catalogue of past lessons; a risk register informed by foresight includes the conditions that could produce new risks. The foresight function supplies the conditions; the risk assessment function assesses the risks those conditions would create. The two functions are distinct, and the handoff between them is the mechanism by which foresight enters the risk register.

Complementarity operates through the scenario set. Scenarios, treated in Chapter 7 of the Planning Manual, describe plausible futures, and each scenario implies a set of risks: the conditions that would arise in that future and the threats those conditions would create. Risk assessment across the scenario set identifies the risks that are robust — present in multiple scenarios — and those that are scenario-specific. The robust risks are priorities for risk treatment; the scenario-specific risks are priorities for monitoring. The scenario set is the bridge between foresight and risk assessment, and its construction is one of the critical outputs of the foresight function.

The institutional arrangements for the two disciplines vary. In some jurisdictions, risk assessment is conducted within the planning authority; in others, a dedicated risk function operates at the centre of government. The United Kingdom's National Risk Register, maintained by the Cabinet Office, incorporates foresight inputs through structured engagement with the Government Office for Science. Singapore's risk assessment is integrated with the strategic planning apparatus in the Prime Minister's Office. What is constant is that the two functions share methods but serve different purposes: risk assessment serves the risk register; foresight serves the assumptions, the scenarios, and the conditions that could produce risks.

3.5 The Boundaries

The boundaries between foresight and its adjacent disciplines matter because each discipline has a distinct purpose, method, and output. When the boundaries are blurred, the disciplines lose their distinctive value. Foresight absorbed into forecasting produces projections without exploration; absorbed into intelligence, it produces threat assessment without long-horizon exploration; absorbed into policy analysis, it produces advocacy without exploration; absorbed into risk assessment, it produces risk lists without the conditions that could create new risks. Each absorption produces a characteristic failure that the institutional arrangements must prevent.

The failure modes are observed in practice across jurisdictions. A foresight unit that drifts into forecasting produces scenarios that converge on a single projection; a unit that drifts into intelligence produces current-threat assessments; a unit that drifts into policy analysis produces recommendations for current decisions. In each case, the range of futures that is foresight's value is lost. The drift is often gradual and is driven by the demands of the customer: the strategy function wants forecasts, the security function wants threat assessments, the policy function wants recommendations. Each demand is legitimate; the failure lies in the foresight function's abandoning its distinct discipline to meet it.

The mature practice maintains the boundaries while collaborating. The foresight function produces foresight products and hands them off to the adjacent functions, which use them for their own purposes. The forecasting function uses foresight to test its baselines; the intelligence function uses foresight to extend its horizon; the policy analysis function uses foresight to frame its options; the risk assessment function uses foresight to expand its risk register. The collaboration is conducted through structured handoffs, not through absorption. Each function maintains its own standards, its own purpose, and its own customer.

Maintaining the boundaries requires institutional discipline. The foresight function's work programme is set by the foresight cycle and the planning cycle, not by the ad hoc demands of

adjacent functions. The foresight function's products are foresight products: scans, scenarios, driver maps, and strategic foresight studies, not forecasts, threat assessments, policy recommendations, or risk registers. The foresight function's staff are trained in foresight methods, not in the methods of adjacent disciplines. The institutional arrangements — reporting lines, work programme governance, product definitions — are the mechanism by which the boundaries are maintained.

3.6 Chapter Summary

Foresight is related to but distinct from forecasting, intelligence analysis, policy analysis, and risk assessment. Each adjacent discipline shares methods with foresight but differs in purpose, horizon, and output. Forecasting supplies the baseline; intelligence analysis assesses current threats; policy analysis assesses options for current decisions; risk assessment identifies and assesses risks. Foresight supplies the conditions against which the baseline is tested, explores the long-term possibilities the intelligence function does not address, examines the futures the policy options must address, and identifies the conditions that could create risks.

The boundaries between the disciplines are the conditions under which each maintains its distinctive value. When foresight is absorbed into an adjacent discipline, it loses the exploration of possible futures that is its contribution. The failure modes are observed in practice: foresight that drifts into forecasting, intelligence, policy analysis, or risk assessment produces the adjacent discipline's output without foresight's distinctive contribution. Each drift is often driven by the demands of the customer and is often gradual.

Maintaining the boundaries while collaborating is the institutional challenge. The foresight function produces foresight products and hands them off to the adjacent functions, which use them for their own purposes. The handoffs are structured, and the institutional arrangements — reporting lines, work programme governance, product definitions — are the mechanism by which the boundaries are sustained. Maintaining the distinction is the condition under which foresight contributes to the plan rather than being absorbed by its adjacent functions.

The foresight practitioner's craft is the conduct of the foresight discipline to a professional standard, distinct from the craft of the forecaster, the intelligence analyst, the policy analyst, and the risk assessor. The methods of foresight — horizon scanning, scenario construction, expert elicitation, driver mapping — are the operational means by which the discipline is practiced. The institutional arrangements are the conditions under which the boundaries are maintained. The boundaries drawn here are the framework within which the practice of foresight is consolidated in the chapters that follow.

Chapter 4 — The Foresight Practitioner

The foresight practitioner is a dedicated professional category.

4.1 The Professional Category

The foresight practitioner is a dedicated professional category within the public service, distinct from the planner and from the intelligence analyst. The defining competence is the conduct of the foresight discipline — horizon scanning, scenario construction, expert elicitation, driver mapping — to a professional standard, under political direction, and within the institutional framework of the state. The category encompasses the foresight analyst, the horizon scanner, the scenario builder, and the foresight facilitator, each bringing a distinct specialism to a common craft. Practitioners serve at every level, from the junior scanner in a sectoral ministry to the head of a central foresight unit; the craft is the same at every level, while the scale and scope differ.

The craft comprises the methods treated in the doctrine and the standards of practice that govern them. Each method has its characteristic outputs, its characteristic failure modes, and its criteria of competent execution. The craft is transmitted through training, through apprenticeship to experienced practitioners, and through participation in the professional community that connects foresight practitioners across jurisdictions. A practitioner who has not mastered the methods has not mastered the craft, however strong their substantive knowledge of the sector under study. Mastery is demonstrated through documented products, peer review, and the practitioner’s capacity to design and lead exercises rather than only contribute to them.

The category has a career path. Practitioners progress from junior scanning and analyst roles to senior roles in which they design and lead foresight exercises, manage teams, and set the work programme of the function. The path is distinct from the planner’s career path and from the intelligence analyst’s career path, though movement between categories over a career is common and often valuable. The path is supported by training institutions, professional networks, and the body of practice that has developed over the past half-century in governmental, corporate, and academic settings. The specific arrangements for rank, certification, and progression are matters for each jurisdiction’s personnel authority.

A professional community transmits the discipline across jurisdictions and across generations of practitioners. The OECD’s strategic foresight community, the European Commission’s foresight networks, the academic foresight community, and the professional associations in many jurisdictions provide the forums through which methods are refined and standards are set. The relationship to the planner is one of complementary crafts: the planner constructs the plan; the foresight practitioner supplies the futures against which the plan is tested. The relationship to the intelligence analyst is one of adjacent disciplines: the intelligence analyst assesses current threats; the foresight practitioner explores future conditions. The boundaries between these categories are maintained by institutional arrangement rather than by personal preference.

4.2 Competencies

The foresight practitioner's competencies comprise horizon scanning, scenario construction, expert elicitation, facilitation, communication, and institutional judgment. Each competency is distinct, with its own methods, its own standards, and its own characteristic failure modes. Each is developed through a combination of formal training, supervised practice, and the practitioner's reflection on their own work. A practitioner strong in some competencies and weak in others is limited in the work they can do; the mature practitioner has command of all six, with depth in two or three. The competencies are treated at the operational level in the chapters that follow; the assessment treats them as elements of a single professional profile.

Horizon scanning is the competency of identifying relevant sources, filtering signals from noise, assessing signals for impact and uncertainty, and maintaining the scan log as institutional memory. It is developed through supervised scanning, through review of the scanner's assessments by experienced practitioners, and through exposure to the full range of source types. Scenario construction is the competency of building internally consistent narratives of plausible futures from the drivers and uncertainties identified through scanning and analysis. It is developed through the construction of scenario sets under supervision, through the critique of those sets by peers and by domain experts, and through reflection on the gap between scenarios constructed and conditions that subsequently emerged. Expert elicitation is the competency of designing and conducting structured elicitation of expert judgment, including the Delphi method, expert panels, and structured workshops; it is developed through training in elicitation methods, through observation of experienced elicitors, and through the practitioner's conduct of elicitations under supervision.

Facilitation is the competency required to run the workshops through which much foresight work is conducted. Scenario workshops, driver mapping sessions, and expert panels all require skilled facilitation: managing the process, ensuring all voices are heard, surfacing dissent, and producing the record. Facilitation is a distinct competency from analysis: the analyst supplies the content and the synthesis; the facilitator supplies the process. Communication is the competency required to translate foresight products into the formats the strategy function and the planning authority can use; the communicator conveys uncertainty without undermining the product's authority and conveys implications without overstating them. Both competencies are developed through supervised practice, through review by experienced practitioners, and through the practitioner's revision of their work in light of how it is received.

Institutional judgment is the competency that integrates the others. It is the ability to assess what foresight is needed, when it is needed, and how it should be presented to be used. The competency includes reading the institutional environment: knowing when an emerging issue is ready for escalation, when a scenario set is robust enough for the strategy function, and when a product should be held for further development. Institutional judgment cannot be taught in the abstract; it is acquired through practice under the guidance of experienced practitioners, through exposure to the decisions the foresight function supports, and through the practitioner's observation of the consequences of their own and others' choices. It is the competency that distinguishes the senior practitioner from the junior, and it is the basis on which the senior practitioner is trusted to design and lead exercises.

4.3 The Duty of Honest Assessment

The foresight practitioner's duty is to produce honest assessments, even when unwelcome. The duty applies to every element of the work: every signal assessed, every scenario constructed, every implication identified, every product disseminated. An assessment that is optimistic because the truth is unwelcome does not meet the standard. A scenario that omits a plausible future because the future is politically uncomfortable fails the duty in the same way, as does a product that softens its implications because the implications are inconvenient. The duty is to the discipline and to the planning function it serves, and it is discharged through honest work regardless of the preferences of the moment.

The pressures against honest assessment are constant. Political pressure may push toward optimism: futures that align with the government's preferences, scenarios that support the government's strategy, implications that confirm the government's decisions. Institutional pressure may push toward conformity: futures that align with the planning authority's assumptions, scenarios that do not require plan revision, implications that do not disturb the existing programme. Customer pressure may push toward usefulness: foresight that produces actionable recommendations rather than uncomfortable exploration, that confirms what the customer already believes rather than challenging it. Each pressure is genuine; the practitioner who denies that the pressures exist is less able to resist them than the practitioner who recognises them.

The integrity required is the capacity to produce honest work despite these pressures. Integrity is sustained as much by institutional arrangement as by personal attribute. The foresight function's independence, its reporting line, its protection from political interference, and its professional standards are the conditions under which integrity is possible. A foresight function that reports to a political office and is subject to political direction in its assessments cannot reliably produce honest work; a foresight function that is institutionally protected can. The practitioner who accepts a position in a function that lacks the protections is accepting a position in which the duty cannot be fully discharged.

The duty extends to the communication of uncertainty. Foresight products convey uncertainty by construction: scenarios describe possible futures, not predicted futures; signals indicate potential change, not certain change. The temptation to convey certainty that the analysis does not support is constant, because certainty is more useful to the customer and more comfortable for the practitioner. The duty is to convey the uncertainty honestly: to state what is known, what is unknown, and what is unknowable, and to let the customer make decisions in light of that. Honest communication of uncertainty is the test by which the duty of honest assessment is operationalised in the product.

4.4 The Relationship to the Planner

The foresight practitioner produces the foresight; the planner uses it. The distinction of roles is fundamental: the practitioner's craft is the exploration of possible futures; the planner's craft is the construction of the plan. The two crafts are complementary, and the handoff between them is one of the critical interfaces in the planning process. The handoff is structured: the practitioner produces defined products; the planner consumes them for defined purposes; the formats, the

timing, and the channels are specified in the planning cycle. A handoff conducted through informal channels, without documentation, fails to discharge the discipline on either side.

The products the practitioner produces include the scan briefing, the emerging issues report, the scenario set, the driver map, and the strategic foresight study. Each product has a defined audience, a defined format, and a defined place in the planning cycle. The planner uses the scan briefing to maintain awareness of emerging change; the emerging issues report to identify conditions that may require plan revision; the scenario set to test the plan's resilience; the driver map to understand the forces shaping the planning environment; the strategic foresight study to inform the plan's diagnosis and strategy. The match between product and use is the substance of the handoff. A mismatch — a product designed for one use, applied to another — degrades both.

The handoff is supported by ongoing collaboration throughout the planning cycle. The practitioner supplies the foresight; the planner identifies the questions the foresight must address; the practitioner revises the products in light of the planner's needs; the planner revises the plan in light of the foresight. The collaboration is conducted through structured interactions: regular briefings, joint workshops, and the planning cycle's formal interfaces. Without the collaboration, the handoff fails: the practitioner produces products the planner cannot use; the planner asks questions the practitioner cannot answer. The collaboration is the means by which the practitioner's craft and the planner's craft are kept aligned as the cycle advances.

The distinction of roles is maintained even in close collaboration. The practitioner does not set the plan's objectives; that is the planner's craft, under political direction. The planner does not construct the scenarios; that is the practitioner's craft. The two roles are distinct in method, in output, and in the decisions they serve. Where the distinction erodes — where the practitioner begins to advocate for objectives, or the planner begins to construct scenarios — both crafts are weakened and the plan's quality suffers.

4.5 The Relationship to Political Authority

The foresight practitioner advises through the planning authority, not directly to political authority, except where institutional arrangements provide direct access. The arrangement is governed by the constitutional position of the civil service and by the institutional form of the foresight function. In the central unit form, the function may have direct access to the head of government; in the distributed model, access is through the planning authority; in the parliamentary committee form, access is through the legislature. Each arrangement has implications for what the practitioner communicates and how. The comparative examples treated in Chapter 2 illustrate the range of arrangements and the conditions under which each operates.

The communication of uncertainty is the practitioner's principal challenge in the relationship to political authority. Political authority wants clarity; foresight produces ranges. Political authority wants recommendations; foresight produces implications. The practitioner's task is to convey the uncertainty honestly while making the implications clear: too much uncertainty and the product is dismissed as inconclusive; too little and the product conveys a false certainty. The judgment required is developed through experience and forms part of the institutional judgment treated in section 4.2.

Where the institutional arrangements provide direct access, the practitioner's responsibility is heightened. Direct access to political authority carries the risk of politicisation: the practitioner may be tempted to tailor the assessment to the political audience, or the political authority may be tempted to direct the assessment. The institutional protections — the reporting line, the professional standards, the documented products — are the conditions under which direct access is safe. Without them, direct access undermines the duty of honest assessment. Singapore's Centre for Strategic Futures within the Prime Minister's Office strategy apparatus, the United Kingdom's Government Office for Science under the Chief Scientific Adviser, and Finland's Committee for the Future as a standing parliamentary committee illustrate the range of arrangements and the conditions under which each operates.

The relationship to political authority is mediated by the planning authority in most arrangements. The planning authority receives the foresight products, integrates them into the plan, and presents the plan to political authority; the practitioner's voice is heard through the plan, not separately. This arrangement has the strength of integration: foresight is used as part of the planning process, rather than as a separate commentary on it. It has the weakness of filtering: the planning authority may moderate the foresight's implications in presenting them to political authority. The mature practice manages the weakness through the documented products: the foresight function's products are on the record, and the planning authority's use of them is traceable.

4.6 Chapter Summary

The foresight practitioner is a dedicated professional category within the public service, distinct from the planner and from the intelligence analyst. The category has a defined craft, a career path, and a professional community. The craft comprises the methods treated in the doctrine; the career path is supported by training institutions and professional networks; the community transmits the discipline across jurisdictions. The relationship to the planner is one of complementary crafts; the relationship to the intelligence analyst is one of adjacent disciplines with distinct objects of assessment.

The practitioner's competencies comprise horizon scanning, scenario construction, expert elicitation, facilitation, communication, and institutional judgment. Each competency is distinct and is developed through training, supervised practice, and reflection. The mature practitioner has command of all six, with depth in two or three. Institutional judgment integrates the others and distinguishes the senior practitioner from the junior.

The duty of honest assessment is the practitioner's obligation to produce honest work even when unwelcome. The pressures against honest assessment are constant: political, institutional, and customer pressures push toward optimism, conformity, and usefulness. The integrity required is sustained as much by institutional arrangement as by personal attribute. The duty includes the communication of uncertainty.

The practitioner's relationship to the planner is one of complementary crafts and structured handoffs, supported by ongoing collaboration throughout the planning cycle. The relationship to political authority is mediated by the planning authority in most arrangements, with direct access

only where institutional protections support it. The communication of uncertainty is the principal challenge in the relationship to political authority, and the judgment it requires is part of the institutional judgment that defines the senior practitioner.

Chapter 5 — The Horizon Scanning Method

Horizon scanning is the systematic collection and assessment of signals that may indicate change.

5.1 What Horizon Scanning Is

Horizon scanning is the systematic, continuous, and structured collection and assessment of signals that may indicate change in the planning environment. A signal is a piece of information — drawn from any source — that, properly interpreted, suggests a trend, a risk, or an opportunity not yet reflected in the plan's baseline. The scan's purpose is to detect such signals, accumulate them over time, and assemble them into a picture of the emerging conditions the plan will have to address. The discipline is the standing sensory capability of the foresight function, maintained continuously regardless of where the planning cycle stands.

The systematic character of the scan distinguishes it from casual reading or unsystematic alertness. Defined sources are scanned on a defined schedule; defined procedures govern collection; defined criteria govern assessment. The methods are documented and applied consistently across scanners and over time, so that a signal assessed in one quarter is comparable to a signal assessed in another. A scan that depends on the individual initiative of the scanner, or that shifts sources and criteria from week to week, does not produce a durable record and does not constitute an institutional capability.

Continuity is the scan's defining demand and its principal resource cost. The scan is conducted on a regular schedule — daily for fast-moving sources, weekly or monthly for slower ones — and is maintained over years, not weeks. A scan that runs intensively during a planning exercise and lapses between exercises loses the temporal depth that allows change to be detected. A single snapshot reveals conditions at a moment; a continuous record reveals the direction and rate of change, and it is the direction and rate of change that the foresight function needs to know.

The scan is distinguished from two adjacent activities that are sometimes confused with it. Ad hoc monitoring is the unsystematic watching of sources in response to events; it is useful in a crisis but does not constitute a capability, because it has no standing character and no accumulated record. Literature review is the systematic survey of published research on a defined question; it is a method used within foresight exercises, but it is bounded by a question and bounded by what is published, whereas the scan is continuous and draws on sources that are not literature — patent filings, demographic data, expert elicitation, regulatory consultations, and field observation. The scan's value lies in its standing character and its breadth, which neither ad hoc monitoring nor literature review provides.

5.2 The Scan Cycle

The scan cycle comprises six stages: collection, filtering, assessment, prioritisation, escalation, and review. The stages are conducted in sequence on each batch of signals and are repeated on a regular schedule that matches the pace of the sources. The cycle is the operational structure of the

scan, and weakness at any stage undermines the whole. A scan with strong collection but weak filtering drowns in noise; a scan with strong filtering but weak assessment produces no judgment; a scan with strong assessment but weak escalation fails to inform the decisions it was built to serve.

Collection is the gathering of signals from the scan's defined source list, on the schedule specified for each source. The collection procedure is documented so that any trained scanner can execute it, and so that the resulting record is comparable across scanners and over time. Filtering follows, applying defined criteria — relevance to the scan's scope, credibility of the source, novelty of the signal — to remove material that should not enter the assessment stage. Filtering is the stage at which the scanner's judgment is most heavily exercised, because the criteria must be applied to a large volume of material in a limited time. The output of filtering is a set of candidate signals that warrant structured assessment.

Assessment is the structured evaluation of each filtered signal against a defined schema: category under the STEEPL framework, potential impact, uncertainty, time horizon, and an assessment narrative that explains what the signal means and why it matters. The assessment is what converts a raw signal into an entry in the scan log, and it is the stage at which the scanner's analytical skill is most directly applied. Prioritisation ranks assessed signals for further investigation or for escalation, applying the criteria of impact, uncertainty, novelty, and relevance to the planning question. Prioritisation prevents the scan from treating all signals as equal and focuses the scanner's attention on the signals most likely to matter to the plan.

Escalation is the transmission of high-priority signals, or of clusters of signals that constitute an emerging issue, to the strategy function and the planning authority. The escalation takes the form of a brief that summarises the signal, the assessment, and the recommended response, and it is the principal channel by which the scan informs decisions. Review is the periodic reassessment of the scan log, the source list, and the scan's scope, conducted to ensure the scan remains current and aligned with the planning function's needs. The cycle is managed by the scanning team leader, runs weekly or monthly depending on the scan's pace, and is documented at each stage so that the scan's operation is visible and auditable.

5.3 The Scan's Scope

The scan's scope is defined by the STEEPL framework: social, technological, economic, environmental, political, and legal. The framework ensures that the scan covers the full range of forces that could shape the planning environment, rather than concentrating on the categories the scanner happens to know best. A scan that covers technology and economy but neglects social and environmental forces misses signals in the uncovered categories and produces a partial picture of the future. The breadth requirement is that all six categories are covered, with the depth of coverage in each adjusted to the planning function's needs and the jurisdiction's circumstances.

The framework is applied through sectoral scans and a cross-sectoral scan. Sectoral scans cover each STEEPL category in depth: the social scanner tracks demographic trends, social attitudes, and cultural change; the technological scanner tracks research and development, technology adoption, and patent activity; the economic scanner tracks growth, trade, and financial conditions; the environmental scanner tracks climate, resource, and ecological change; the political

scanner tracks electoral trends, governance change, and political instability; the legal scanner tracks regulatory change, judicial decisions, and international law. The cross-sectoral scan covers the intersections among categories — automation and labour markets, climate and migration, demographic change and fiscal sustainability — where the most consequential signals often arise.

The breadth requirement is demanding and is rarely met by a single team. Covering six categories in depth requires either a large central scanning team or a distributed model in which sectoral scanners are located in sectoral ministries and coordinated by a central foresight unit. The distributed model is the common solution in jurisdictions of any scale: the central unit maintains the cross-sectoral scan, the scan log, and the standards of method, while sectoral scanners in ministries maintain the sectoral scans under those standards. The model requires shared assessment criteria, a common scan log, and regular coordination meetings, without which the sectoral scans produce inconsistent records that cannot be combined into a coherent picture.

The scope is reviewed periodically to ensure it remains aligned with the planning function's needs. A scan that continues to cover categories no longer relevant to the plan wastes effort; a scan that misses categories newly relevant leaves gaps that may prove costly. The review is conducted annually or at the start of each planning cycle, and it adjusts the source list, the sectoral assignments, and the depth of coverage in each category. The review is also the occasion at which the scan's performance is assessed: which signals detected by the scan proved significant, which significant developments the scan missed, and what changes to method or sources would improve the scan's contribution. The review is the mechanism by which the scan learns from its own record.

5.4 The Scan's Output

The scan produces three standing outputs: the scan log, the scan briefing, and the emerging issues report. Each has a defined audience, a defined format, and a defined schedule, and each is produced regardless of whether a foresight exercise is under way. The outputs are the scan's contribution to the foresight function and to the planning cycle, and their quality determines the scan's value to the institution. A scan that produces outputs on the planning cycle's calendar but of poor quality delivers little of value on time; a scan that produces high-quality outputs on a schedule unconnected to the planning cycle delivers them when they cannot be used.

The scan log is the durable record of the scanning function. It contains every assessed signal, with its category, potential impact, uncertainty, time horizon, assessment narrative, source, and date, and it accumulates over years into the scan's institutional memory. The log's structure, accessibility, and maintenance are the conditions under which the scan's memory is preserved across staff turnover and across changes of government. The scan briefing is the regular summary of the scan's recent findings, produced weekly or monthly for the foresight function, the planning authority, and other standing audiences. The briefing highlights new signals, reassessed signals, and signals that have been escalated, and it is the scan's principal operational output.

The emerging issues report is the periodic synthesis of the scan's findings into identified emerging issues. An emerging issue is a cluster of related signals that, taken together, suggest a development the plan should address — a development that is not yet a trend but that, if it consolidates, would have significant implications. The report is produced quarterly, semi-annually,

or annually depending on the scan's pace and the planning cycle's needs, and it sets out each emerging issue with its supporting signals, its potential impact, its time horizon, and its implications for the plan. The report is the scan's principal contribution to the strategy function and to the planning cycle's diagnosis phase.

The outputs' relationship to the strategy function and the planning cycle is structured. The scan briefing informs the strategy function's ongoing awareness of emerging change and supplies the material for the strategic conversation that the strategy function maintains between formal exercises. The emerging issues report informs the planning cycle's diagnosis and strategy phases, supplying the candidate issues the plan must address and the evidence base for the issues it selects. The scan log informs the foresight exercises conducted during the cycle, supplying the raw material for scenario construction, driver mapping, and expert elicitation. The outputs are produced on the planning cycle's calendar, with the briefing arriving in time for the strategy conversation and the emerging issues report arriving in time for the diagnosis phase.

5.5 Institutionalising the Scan

The scan is a continuous function, conducted on a standing schedule and maintained across the planning cycle's phases. Institutionalising it requires four elements: dedicated staff, a defined budget, a documented method, and an institutional memory. Each is necessary, and the absence of any one causes the scan to lapse. A scan staffed by people whose primary duties lie elsewhere is neglected when those duties press; a scan without a budget cannot maintain its sources or its tools; a scan without a documented method depends on the individuals who conduct it and dies when they leave; a scan without an institutional memory loses its continuity at every staff transition. The four elements together are the conditions under which the scan survives the ordinary pressures of government and continues across changes of personnel and of political direction.

The dedicated scanning team is the core of the institutional scan. The team comprises the scanners who conduct collection, filtering, and assessment, and the team leader who manages the cycle, produces the briefings, and represents the scan to the strategy function. The team may be small — two or three scanners in a small jurisdiction — or substantial — a dozen or more in a large jurisdiction operating under a distributed model. The team's size is governed by the scan's scope and the jurisdiction's scale; what is constant is that scanning is the team's primary duty, and the arrangement fails when scanning is treated as an additional task assigned to officials with other responsibilities. Singapore's Centre for Strategic Futures and the United Kingdom's Government Office for Science maintain dedicated scanning capacity of this kind, and the practice is the standard against which other arrangements are measured.

The scan's budget covers three categories of expenditure: sources, tools, and travel. Sources include subscriptions to the technical and scientific press, databases of patent filings and scientific literature, and access to expert networks; tools include the scan log software, the source management system, and the analytical tools that support assessment and prioritisation; travel covers attendance at conferences and professional meetings where scanners maintain their networks and detect signals that have not yet reached published sources. A scan without an adequate budget degrades as sources are dropped, tools become obsolete, and the scanners lose

contact with the professional community that keeps their craft current. The budget is reviewed annually, alongside the scope review, to ensure it remains sufficient to the scan's task.

The scan's institutional memory is the scan log, and the log's quality determines whether the scan accumulates knowledge or merely repeats itself. The log must be structured so that signals can be retrieved by category, by source, by date, and by emerging issue; it must be accessible to the scanners and to the analysts who draw on it; and it must be maintained so that reassessments and disconfirmations are recorded alongside the original entries. A scan log that is poorly structured, inaccessible, or unmaintained loses its value as memory and reduces the scan to a sequence of snapshots rather than a continuous record. The institutional memory is the scan's principal asset and its principal obligation to maintain, because the log is what survives staff turnover and what allows the scan to learn from its own past.

5.6 Chapter Summary

Horizon scanning is the systematic, continuous, and structured collection and assessment of signals that may indicate change in the planning environment. It is the standing sensory capability of the foresight function, maintained on a regular schedule and over years. The scan is distinguished from ad hoc monitoring by its systematic and continuous character, and from literature review by its breadth of sources and its unbounded scope.

The scan cycle comprises six stages — collection, filtering, assessment, prioritisation, escalation, and review — conducted in sequence on each batch of signals and repeated on a regular schedule. The scan's scope is defined by the STEEPL framework, ensuring coverage of social, technological, economic, environmental, political, and legal forces. The scope is implemented through sectoral scans and a cross-sectoral scan, typically coordinated by a central foresight unit with sectoral scanners located in ministries.

The scan produces three standing outputs: the scan log as durable record, the scan briefing as regular operational summary, and the emerging issues report as periodic synthesis. The outputs are produced on the planning cycle's calendar and inform the strategy function and the planning cycle's diagnosis and strategy phases. Institutionalising the scan requires dedicated staff, a defined budget, a documented method, and an institutional memory; the absence of any one causes the scan to lapse. The chapters that follow treat the sources, the assessment methodology, and the emerging issues process in operational detail.

Chapter 6 — Sources and Source Assessment

Horizon scanning draws on a range of sources.

6.1 Source Types

Horizon scanning draws on nine principal source types: the scientific literature, the technical press, government reports from other jurisdictions, international organisation assessments, patent filings, demographic data, expert elicitation, social media, and grey literature. Each type has characteristic strengths and characteristic weaknesses, and a defensible scan draws on all of them. A scan that relies on one or two source types misses signals that the omitted types would reveal and is exposed to whatever biases the included sources carry. Source selection is therefore a matter of portfolio construction rather than preference, and the portfolio is documented in the scan's source list.

The scientific literature comprises peer-reviewed journals and conference proceedings across all fields. Its strength is rigour: papers undergo review, methods are documented, and findings are subject to replication. Its weaknesses are lag, since the interval between research and publication can run to years, and narrowness, since papers address specific questions rather than broad trends. The technical press covers research, development, industry, and policy for a professional audience and offers currency that the scientific literature cannot match. Its quality ranges from rigorous outlets such as Nature, Science, and the IEEE publications to trade publications that prize novelty over accuracy.

Government reports from other jurisdictions provide comparative material: a development in one jurisdiction often signals a development that will affect others. The publications of the UK Government Office for Science, Singapore's Centre for Strategic Futures, the Korean KISTEP, and Finland's SITRA are widely read across jurisdictions as reference material. International organisation assessments — from the OECD, the World Bank, the International Energy Agency, the United Nations, and similar bodies — synthesise material across jurisdictions and provide a baseline against which to assess local conditions. Patent filings reveal technological development before it reaches the market, since filings become public record well before products appear; the databases of WIPO, the EPO, and the USPTO are the standard sources. Demographic data, drawn from the national statistics office and from international compilations, provides the quantitative foundation for social and economic assessment.

Expert elicitation, treated in detail in section 6.5, provides access to judgment that the published record does not yet contain. Social media gives early indication of public attitudes and emerging concerns, though with significant quality and bias challenges that demand careful handling. Grey literature — reports by think tanks, research institutes, and consulting firms — provides analysis that may not have reached the scientific literature or the technical press. Each source type is assessed for quality and bias before its signals enter the scan log, and the assessment is documented alongside the signal. The portfolio is reviewed when the scan's priorities change, so that source selection tracks the questions the scan is asked to answer.

6.2 Source Quality

Source quality is the assessed fitness of a source for use in the scan, judged against five criteria: authority, accuracy, currency, coverage, and methodology. Authority is the standing of the source in its field: a peer-reviewed journal carries higher authority than a blog, and a report by an established research institute carries higher authority than a consulting firm's marketing material. Accuracy is the source's track record of being correct, established through observation over time. Currency is the recency of the source's information, since a signal from a current report carries more weight than one from a report several years old. The three criteria are assessed together with coverage and methodology in the next step of the framework.

Coverage is the breadth of the source's treatment: a source that addresses the question in depth is more valuable than one that mentions it in passing. Methodology is the transparency of the source's method, since a report that documents its data, its assumptions, and its analytical steps is more valuable than one that asserts conclusions without basis. The five criteria are applied together rather than individually, and a source is assessed as strong or weak across the set. A source strong on authority but weak on currency may be valuable for context but not for emerging signals. A source strong on currency but weak on methodology may be useful as a lead but not as a basis for assessment.

The quality framework is documented and applied consistently across the scan. Each source in the scan's source list carries a quality assessment, recorded in the source management system. The assessment is reviewed periodically, because sources change: journals decline, think tanks shift focus, technical publications change editors, and new sources emerge. A source list that is not reviewed degrades over time, as strong sources weaken and weak sources are not culled. The review cycle is therefore part of the framework, integrated into the scan's standard procedures rather than treated as an exceptional exercise.

Documentation of source quality in the scan log is the condition under which the scan's assessments are defensible. Each signal in the log is linked to its source, and the source's quality assessment is part of the record. A signal from a high-authority source is treated differently from a signal from a low-authority source, and the difference is visible in the log. The documentation allows the scan's users — the strategy function, the planning authority, the foresight analysts — to assess the basis of the scan's findings and to weight them accordingly. Without the documentation, the scan's assessments are assertions; with it, they are evidence-based judgments.

6.3 Source Bias

Source bias is the systematic distortion of a source's output by factors other than the evidence. Four principal biases affect foresight sources: ideological bias, commercial bias, institutional bias, and publication bias. Each operates differently, and each must be assessed and managed. A scan that does not assess bias is exposed to it: signals from biased sources enter the log without correction, and the scan's findings reflect the biases of its sources rather than the underlying conditions. Bias assessment is therefore conducted for every source in the source list and recorded alongside the quality assessment.

Ideological bias is the distortion of a source's output by its political or ideological commitments. Think tanks, advocacy organisations, and some media outlets hold ideological positions that shape their reporting. The bias operates through topic selection, question framing, and evidence interpretation rather than through overt assertion, which makes it harder to detect. Commercial bias is the distortion of a source's output by its commercial interests. Consulting firms, industry associations, and some technical publications have commercial interests that shape their analysis: a report by a consulting firm that sells services in the field it assesses is commercially biased, however rigorous its method appears.

Institutional bias is the distortion of a source's output by its institutional position. Government reports, international organisation assessments, and research institute publications reflect the positions of the institutions that produce them. A government report may understate problems that reflect poorly on the government; an international organisation assessment may reflect the consensus among its members rather than the underlying evidence. Publication bias is the distortion of the published record by the factors that determine what is published. Studies with positive findings are more likely to be published than studies with negative findings, and studies that support the prevailing view are more likely to be published than studies that challenge it.

Management of bias is through source triangulation. A signal that appears in only one source, or only in sources that share a bias, is treated with caution. A signal that appears in multiple independent sources, with different biases, is treated with more confidence. Triangulation is the scanner's principal defence against bias: by seeking confirmation across sources with different biases, the scanner distinguishes signals that reflect underlying conditions from signals that reflect source bias. Triangulation is documented in the scan log, where each signal records the sources it appears in, the biases of those sources, and the scanner's assessment of the signal's reliability in light of them.

6.4 The Breadth-versus-Depth Trade-off

The breadth-versus-depth trade-off is the allocation of finite scanning effort between covering many sources and covering each source thoroughly. The scan's scope is broad — six STEEPL categories, each with many sources — and the scanner's time is finite. A scan that prioritises breadth covers many sources but loses depth in each; a scan that prioritises depth covers few sources and misses signals in the uncovered ones. Neither extreme is satisfactory, and the trade-off must be managed deliberately through a framework that allocates depth where it is most valuable. The framework is the scanner's principal response to the resource constraint under which all scanning is conducted.

The framework rests on a tiered source list. A small set of core sources is scanned in depth, a larger set of secondary sources is scanned at the headline level, and a long list of monitoring sources is checked periodically for relevant items. Core sources are those whose coverage is most directly relevant to the scan's priority questions; secondary sources provide supporting coverage; monitoring sources provide a safety net for developments outside the scan's primary focus. The tiering ensures depth where it is most valuable and breadth where depth is not affordable. Tier assignments are documented in the scan's procedures and reviewed when priorities change.

Scanning depth is adjusted to the source's tier. Core sources are read in full; secondary sources are scanned at the headline and abstract level; monitoring sources are checked by keyword or alert. Assessment depth is adjusted in parallel to the signal's priority. High-priority signals are assessed in depth: the scanner investigates the source, seeks triangulation, assesses the implications, and produces a full assessment narrative. Low-priority signals are recorded briefly with source and note, and the adjustment rules concentrate the scanner's effort where it is most valuable.

The distributed model, treated in Chapter 2, is the principal means by which a broad and deep scan is sustained within finite resources. Sectoral scanners in ministries cover their sectors in depth, and the central foresight unit covers the cross-sectoral scan and the integration. The distribution allows depth across all sectors without requiring each scanner to cover everything. Coordination is the condition of the model's success: shared standards, a common scan log, and regular communication among scanners. Without coordination, the distributed model produces inconsistent depth and gaps at the intersections.

6.5 Expert Elicitation as a Source

Expert elicitation is a structured source that provides access to judgment not yet in the published record. It comprises three principal methods: the Delphi method, expert panels, and structured workshops. Each method has its own design, strengths, and weaknesses. Expert elicitation is used when the published record is insufficient — when the question is too new for the literature, too specific for the technical press, or too uncertain for quantitative methods. It is a complement to published sources, used where they do not reach, and the methods and experts consulted are documented in the scan log.

The Delphi method is a structured elicitation conducted over multiple rounds. A panel of experts is surveyed on a question, the responses are summarised and fed back to the panel, the panel is re-surveyed in light of the feedback, and the process is repeated until the responses converge or the divergence is clearly identified. The Delphi's strength is the reduction of dominance: experts respond independently, and the aggregation prevents any single expert from dominating. Its weakness is the time it takes and the difficulty of sustaining expert engagement across multiple rounds. The method is well suited to questions where the expert pool is dispersed and where independent judgment matters more than debate.

Expert panels bring experts together for structured discussion. The panel's design is critical: the question must be clearly framed, the experts must be selected for diversity of perspective, and the discussion must be facilitated to surface dissent rather than suppress it. The panel's strength is the depth of discussion and the surfacing of reasoning; its weakness is the risk of dominance by the most assertive expert and the pressure toward consensus. Structured workshops extend the panel format to include non-experts — planners, stakeholders, decision-makers — and are used to explore implications as well as to elicit judgment. Workshops are particularly useful when the question concerns the implications of a development for a specific institution or sector.

Management of expert bias is built into the design of any expert elicitation. Experts carry disciplinary biases that frame questions in their field's terms, institutional biases that reflect their employer's position, and personal biases that reflect their preferences. The design addresses bias

through diverse selection of experts across discipline, institution, and perspective, through structuring that secures independent response before discussion where possible, and through documentation that records disagreement as well as consensus. Synthesis of expert input is the scanner's responsibility: the input is recorded, areas of agreement and disagreement are identified, and the synthesis is documented in the scan log with the experts' reasoning and the scanner's assessment. Expert elicitation is treated in more detail in Part IV.

6.6 Chapter Summary

Horizon scanning draws on nine principal source types: the scientific literature, the technical press, government reports from other jurisdictions, international organisation assessments, patent filings, demographic data, expert elicitation, social media, and grey literature. Each type has characteristic strengths and weaknesses, and the scan draws on all of them to avoid the biases of relying on one or two. Source selection is portfolio construction, and the portfolio is documented in the scan's source list.

Source quality is assessed against five criteria: authority, accuracy, currency, coverage, and methodology. The quality framework is documented and applied consistently, and each source's quality assessment is recorded in the source management system. Documentation of source quality in the scan log is the condition under which the scan's assessments are defensible.

Source bias is assessed across four dimensions: ideological, commercial, institutional, and publication bias. Management of bias is through source triangulation: signals confirmed across sources with different biases carry more confidence than signals from a single source or from sources that share a bias. Triangulation is documented in the scan log alongside the signal.

The breadth-versus-depth trade-off is managed through a tiered source list, adjusted scanning depth, and adjusted assessment depth. The distributed model allows depth across all sectors within finite resources, with coordination by the central foresight unit. Expert elicitation provides access to judgment not yet in the published record through the Delphi method, expert panels, and structured workshops, with bias addressed in the design of the elicitation.

Chapter 7 — Weak Signals and Emerging Issues

Weak signals are the particular object of horizon scanning.

7.1 What Weak Signals Are

A weak signal is an early indicator of a change that, if it develops, would have significant implications for the plan. The defining attributes are three: novelty, potential impact, and uncertainty. Novelty means the signal has not yet consolidated into a recognised trend; it appears in isolated observations rather than in established data series. Potential impact means that, if the underlying change develops, it would materially affect the planning assumptions, the plan's commitments, or the conditions of its success. Uncertainty means the signal may consolidate into a trend, may dissipate, or may evolve in an unanticipated direction.

The distinction between a weak signal and a trend is one of consolidation. A trend is supported by sufficient evidence to support projection; a weak signal is supported by insufficient evidence to support projection but sufficient to warrant attention. The scanner logs the weak signal because the cost of missing a signal that later consolidates exceeds the cost of logging one that later dissipates. The discipline requires the scanner to maintain a record of signals that, if consolidated, would matter; the discipline does not require the scanner to predict which signals will consolidate. The record becomes the input to the emerging issues methodology.

Weak signals are observed across the same domains as horizon scanning: social, technological, economic, environmental, political, and legal. A weak signal may be a single observation in a technical publication, a pattern noted by an expert in an unrelated field, an anomaly in an early data series, or a regulatory change in another jurisdiction that has not yet been considered locally. The Singapore Centre for Strategic Futures and the United Kingdom Government Office for Science both maintain weak-signal registers that capture observations across these domains. The mature practice does not restrict the source of weak signals; it restricts the assessment of them.

7.2 The Distinction from Noise

The scanner's challenge is to distinguish weak signals from the noise of irrelevant information. Noise is the volume of material collected during horizon scanning that does not constitute a signal. The scan log will contain more noise than signal; this is intrinsic to the discipline and is not a sign of poor practice. The challenge is structural: weak signals are by definition rare, and the methods that detect them must accept a high volume of material to find the few that warrant attention. A scan that excludes noise at the point of collection excludes weak signals as well, because the difference between them is often not visible until assessment.

The framework for distinguishing signal from noise rests on three filters. The first is source quality: the scanner assesses whether the source is reliable, whether the observation is verifiable, and whether the source has a record of early identification. The second is relevance: the scanner

assesses whether the observation, if it developed, would bear on the planning question, the assumptions, or the commitments. The third is corroboration: the scanner looks for the same observation in independent sources, recognising that a single observation may be signal or noise but that two or three independent observations of the same phenomenon raise the probability of signal. The filters are applied sequentially; an observation that fails source quality is not assessed for relevance, and one that fails relevance is not assessed for corroboration.

The role of judgment is central and irreducible. The framework narrows the field; the final assessment rests with the scanner, who must decide whether the observation, taken together with its source and corroboration, warrants a log entry and an escalation flag. Different scanners will reach different judgments on the same observation; this is intrinsic to the discipline and is managed rather than eliminated. Established practice manages it through peer review of the scan log, through structured discussion of contested entries, and through periodic external review by experts from outside the scanning team. The Singapore and Helsinki foresight functions both operate peer-review cycles; the Finnish Committee for the Future provides a parliamentary review channel for signals that carry political as well as analytical implications.

The scan log is the institutional record through which the framework operates. Each entry records the signal, the source, the date of collection, the scanner's assessment under the three filters, and the disposition: retained for monitoring, escalated, or closed. The log is reviewed on a defined schedule, typically weekly for new entries and quarterly for retained entries, so that signals do not accumulate without assessment. The structure of the log enables the clustering that produces emerging issues; without a structured log, the scanner would have observations but no methodology. The discipline treats the log as the primary output of horizon scanning, with individual assessments and reports as derived and subsidiary products.

7.3 The Emerging Issues Methodology

An emerging issue is a cluster of weak signals that, taken together, suggests a change that is beginning to consolidate. The relationship between a weak signal and an emerging issue is one of accumulation: a weak signal is a single observation; an emerging issue is a pattern of observations sufficient to warrant attention from the strategy function. The methodology identifies emerging issues by clustering weak signals, assessing the clusters, and naming the issue that each cluster represents. The output is a small number of emerging issues, each supported by a documented cluster of weak signals and an assessment of its potential implications.

Clustering groups weak signals that bear on the same phenomenon. The scanner reviews the scan log on a defined schedule, typically quarterly or semi-annually, and groups entries that address the same change, the same driver, or the same sector. The grouping is conducted by the scanning team in a structured workshop; the meaning of similarity is a judgment that requires interpretation. A cluster may bring together signals from different domains: a technological signal, a regulatory signal from another jurisdiction, and an expert observation may all bear on the same emerging issue. The output of clustering is a set of candidate clusters, each with the signals it contains and a preliminary statement of the issue it may represent.

Each candidate cluster is then assessed for three characteristics: the strength of the cluster, the potential impact of the issue it represents, and the horizon over which the issue may develop. Strength is a function of the number of signals, the independence of their sources, and the consistency of their direction. Potential impact is assessed against the planning question, the assumptions, and the commitments. Horizon is the period over which the issue may move from emergence to materiality. A cluster that is strong, high-impact, and short-horizon is escalated immediately, while a cluster that is weak, low-impact, or long-horizon is held for further monitoring; the assessment is documented, and the cluster is either promoted to an emerging issue, retained as a candidate, or closed.

The relationship to the scan log is continuous and bidirectional. The scan log feeds the clustering; the clustering feeds back into the scan log by directing the scanner's attention to signals that bear on active clusters. A new signal that fits an existing cluster is logged against it; a new signal that does not fit any cluster may initiate a new one. The scan log also retains the closed clusters, which form the historical record against which future clusters are assessed. The methodology produces a current assessment that is revised each cycle, with the revisions documented so that the reasoning can be reviewed.

7.4 The Escalation Process

Escalation is the process by which a weak signal or an emerging issue is brought to the attention of the strategy function. The scanner does not act on weak signals directly; the scanner's role is to detect, assess, and escalate. The strategy function decides whether and how to act. Escalation is therefore a handover, governed by criteria that distinguish the signals and issues that warrant it from those that warrant continued monitoring. The criteria are documented in advance and applied consistently, so that the strategy function receives a predictable flow of escalations rather than an ad hoc stream.

The criteria for escalation rest on three thresholds. The first is materiality: the issue, if it developed, would require a change to the plan's assumptions, a revision to a commitment, or a new strategic risk entry. The second is plausibility: the cluster is strong enough that the issue cannot reasonably be dismissed as noise. The third is timeliness: the horizon over which the issue may materialise is short enough that acting now is preferable to waiting for further consolidation. An issue that meets all three thresholds is escalated, while one that meets one or two is documented and held for the next assessment cycle; the thresholds are set in consultation with the strategy function so that the volume and nature of escalations match the function's capacity to consider them.

The escalation brief is a short document, typically two to four pages, that presents the issue, the supporting cluster, the assessment, and the proposed response. The brief states the issue in a single paragraph; lists the weak signals that constitute the cluster, with sources and dates; assesses materiality, plausibility, and timeliness; and proposes either a strategic response, a deeper analysis, or continued monitoring with specified indicators. The brief does not recommend a preferred course of action; that is the strategy function's role. The brief presents the evidence and the assessment, and leaves the choice to the function that bears the choice.

The follow-up completes the cycle. The strategy function responds to the brief with a decision: to act, to commission further analysis, to monitor with specified indicators, or to set aside. The decision is recorded against the escalation in the scan log, so that the disposition of every escalated issue is traceable. Where the decision is to monitor, the scanner adds the specified indicators to the routine scan and reports on them at the next assessment cycle; where the decision is to set aside, the scanner retains the cluster in the historical record and flags it for review if new signals emerge. The follow-up ensures that the escalation produces a recorded response and a defined next step.

7.5 Worked Example: Detecting a Weak Signal

Consider a planning authority preparing a five-year plan for tertiary education. The scanner, conducting routine horizon scanning, encounters a working paper from a research institute in another jurisdiction reporting a sharp increase in enrolment in short-cycle vocational credentials among workers aged 35 to 50, alongside a decline in enrolment in traditional master's degrees in the same cohort. The observation is novel; the scanner has not previously encountered a report of mid-career learners shifting from degree programmes to short-cycle credentials at this scale. The observation is potentially material: if the pattern developed locally, it would affect demand for the authority's tertiary institutions and the assumptions about credential structures in the plan. The scanner logs the signal, attaches the source, and flags it for corroboration.

At the next scan cycle, the scanner identifies two further signals that bear on the same phenomenon. The first is a regulatory change in a second jurisdiction that formally recognises short-cycle credentials for civil service recruitment, removing a barrier that previously favoured degree holders. The second is an expert observation, recorded in a workshop, that employers in a third jurisdiction are beginning to advertise roles against specific skills rather than against degree qualifications. The three signals are independent in source, consistent in direction, and together suggest a shift in the relationship between credentials and employment. The scanner clusters the three signals, drafts a preliminary statement of the emerging issue, and presents the cluster at the next clustering workshop.

The workshop assesses the cluster. Strength is moderate: three independent signals, consistent in direction, but all from other jurisdictions and none yet supported by local data. Potential impact is high: if the shift developed locally, the plan's assumptions about demand for master's-level provision would require revision, and the credential architecture would need to accommodate short-cycle programmes. Horizon is uncertain: the signals are from comparable jurisdictions and could appear locally within the plan period, or could fail to appear at all. The workshop concludes that the cluster meets the plausibility and materiality thresholds for escalation but not the timeliness threshold, because the horizon is too uncertain; the cluster is promoted to an emerging issue, retained for monitoring, and assigned indicators: local employer surveys, enrolment data disaggregated by age and credential type, and a watch on regulatory developments in comparable jurisdictions.

Six months later, the local employer survey reports a measurable shift in role specifications away from degree requirements and toward skill specifications. The signal corroborates the cluster, the horizon shortens, and the scanner prepares an escalation brief. The brief presents the issue, the

four supporting signals, the assessment of materiality, plausibility, and timeliness, and proposes that the strategy function commission a deeper analysis of the credential architecture and consider a pilot of short-cycle provision. The strategy function responds by commissioning the analysis and adding the indicators to the routine monitoring system. The cluster is retained in the scan log, marked as escalated, and scheduled for review at the next cycle.

7.6 Chapter Summary

A weak signal is an early indicator of a change that, if it developed, would have significant implications for the plan. The defining attributes are novelty, potential impact, and uncertainty. The scanner's challenge is to distinguish weak signals from noise, using a framework of source quality, relevance, and corroboration, with judgment as the irreducible final filter. An emerging issue is a cluster of weak signals that, taken together, suggests a change beginning to consolidate; the methodology clusters, assesses, and names emerging issues on a defined cycle, with the scan log feeding the clustering and the clustering feeding back into the log.

Escalation is the handover of a weak signal or emerging issue to the strategy function, governed by criteria of materiality, plausibility, and timeliness. The escalation brief presents the issue, the supporting cluster, the assessment, and a proposed response; the follow-up records the strategy function's decision and defines the next step. The worked example illustrates the cycle from detection through clustering, assessment, escalation, and follow-up, and shows how the methodology operates when a signal that initially fails the timeliness threshold later meets it through corroboration.

Chapter 8 — Signal Assessment and the Scan Log

The scan log is the foresight function's institutional memory.

8.1 The Structured Assessment

The structured assessment of each signal is the stage at which the signal becomes an entry in the scan log. The assessment applies a defined framework: each signal is assessed for category, potential impact, uncertainty, time horizon, and assessment narrative. The framework ensures that every signal in the log is assessed to the same standard and that the log's entries are comparable across scanners and over time. A log in which signals are assessed by different criteria, or by impression instead of by framework, fails as a structured record and cannot support the scan's downstream uses. The assessment is conducted by the scanner who detected the signal, with the team leader reviewing a sample of entries to confirm that the framework is applied consistently across scanners and over time.

The category is the STEEPL classification of the signal: social, technological, economic, environmental, political, or legal. Some signals cross categories and are classified by their primary category with secondary categories noted, so that the log preserves the cross-sectoral character of the change instead of forcing a single classification. The category allows the log to be searched by sector and allows the scan's outputs to be organised by sector for the use of sectoral ministries and policy units. The potential impact is the magnitude of the implications if the signal develops, assessed on a defined scale of high, medium, or low, with criteria set for each level. The assessment is the scanner's judgment, informed by the signal's content, the scanner's knowledge of the sector, and the scanner's understanding of the planning question the scan serves.

The uncertainty is the degree to which the signal's implications are unclear. Uncertainty is assessed on a defined scale — high, medium, low — with criteria for each level that distinguish the source of uncertainty: sparse data, conflicting sources, or genuine indeterminacy in the underlying process. A signal with high uncertainty may be significant, because uncertainty often reflects novelty; a signal with low uncertainty may be less significant, because the implications are already clear. The time horizon is the period over which the signal's implications would arise: near-term (within the plan horizon), medium-term (within ten years), long-term (beyond ten years). The time horizon allows the log to be searched by period and allows the scan's outputs to be organised by horizon, matching the different horizons of the plan, the medium-term fiscal framework, and the long-term infrastructure cycle.

The assessment narrative is the scanner's written account of the signal, its source, its implications, and the scanner's reasoning. The narrative is the substantive content of the log entry, and it is the basis on which the signal will be assessed by the scan's users — the team leader, the foresight analysts who draw on the log for exercises, and the strategy function to which signals may be escalated. The narrative is concise — a paragraph or two — and is written to be read by someone who is not the scanner, with the source's quality and bias stated, the scanner's assessment of the signal's reliability recorded, and the scanner's judgment about the signal's significance

explained. The narrative is the entry that makes the log more than a register of metadata; without it, the category, impact, uncertainty, and time horizon fields cannot be interpreted by anyone who was not present at the assessment. The assessment framework, applied consistently across all entries, is the condition under which the scan log is a structured record rather than a collection of notes.

8.2 Prioritisation

Prioritisation ranks assessed signals for further investigation or escalation. The prioritisation applies four criteria: impact, uncertainty, novelty, and relevance to the planning question. The criteria are combined into a prioritisation matrix that allows signals to be ranked and compared across the log. The matrix is the scanner's principal tool for managing attention: with limited resources, the scan must concentrate on the signals that most warrant it, and the matrix is the framework for that concentration. Prioritisation is revisited each time new signals enter the log, so that the ranking reflects the current state of the scan rather than a snapshot from an earlier date.

The prioritisation matrix plots signals on two axes: impact and uncertainty. High-impact, high-uncertainty signals are the highest priority: they have substantial implications and their development is unclear, so they warrant the most attention. High-impact, low-uncertainty signals are the next priority: they have substantial implications and their development is clear, so they warrant attention to the implications rather than to the signal itself. Low-impact signals are lower priority regardless of uncertainty, because their implications are limited. The matrix is supplemented by the novelty and relevance criteria: a novel signal is prioritised over a familiar one at the same impact and uncertainty; a signal relevant to the current planning question is prioritised over one relevant to a future cycle.

The prioritisation determines the scanner's allocation of effort. High-priority signals are investigated in depth: the scanner seeks triangulation from additional sources, assesses the implications for the plan's assumptions and risks, and prepares the signal for escalation if warranted. Medium-priority signals are monitored: the scanner records them and checks for developments at the regular review, with the depth of investigation deferred until the signal rises in priority. Low-priority signals are recorded and held: they enter the log but do not receive active attention unless new information raises their priority. The allocation is documented in the log, so the scanner's effort is traceable and the prioritisation can be reviewed by the team leader and audited over time.

The prioritisation matrix is reviewed periodically to ensure the criteria and their application remain appropriate. The review considers whether the matrix is producing the right priorities: are the escalated signals the ones that warrant escalation, and are the held signals being reassessed when new information arrives? The review is conducted by the scanning team leader and is the mechanism by which the prioritisation is refined over time, as the planning question evolves and as the scan's understanding of the signal landscape matures. A matrix that is not reviewed drifts: the criteria that were appropriate when the scan was established may no longer match the planning function's needs, and the prioritisation loses its alignment with the decisions the scan is meant to inform. A matrix that is reviewed remains aligned with the scan's purpose and the planning

function's needs, and the discipline of the review is what keeps the matrix producing useful priorities over time.

8.3 The Scan Log as Institutional Memory

The scan log is the durable record of the scanning function and the foresight function's institutional memory. It contains every assessed signal, with its category, impact, uncertainty, time horizon, assessment narrative, source, and date, together with the links to related signals and to any emerging issue the signal has been clustered into. The log preserves the scan's work across personnel turnover, across planning cycles, and across changes of government, and it is the record from which the scan's history can be reconstructed. The log is also the source from which the foresight function's products — the scan briefing, the emerging issues report, the scenario exercise — are constructed, so that each product can be traced to the signals on which it rests. Without the log, the scan's work is lost each time a scanner leaves or a cycle ends; with the log, the work accumulates and the scan's memory grows.

The log's structure determines whether it can function as memory. The log is organised to be searchable: by category, by date, by source, by priority, by emerging issue, and by the planning question to which the signal is relevant. Each entry is linked to its source, to related signals, and to any emerging issue it has been clustered into, so that the log forms a network of references instead of a flat list. The structure allows a new scanner to understand the log's contents quickly, to find signals relevant to a current question, and to trace the history of an emerging issue from its first detection through its escalation. A log that lacks structure — a collection of files or notes without organisation — becomes an archive that cannot be used as memory.

The log's role in preserving knowledge across personnel turnover is its principal institutional function. Scanners move, retire, and transfer; the log remains, and the work that the scan has accumulated is not lost with the scanner who produced it. A new scanner joining the team can review the log to understand what has been scanned, what has been assessed, and what has been escalated, and can build on the work of predecessors instead of starting afresh. The foresight bodies that have sustained scanning over decades — Singapore's Centre for Strategic Futures, the United Kingdom's Government Office for Science, and the secretariats supporting Finland's Committee for the Future — have done so in part because their scan records persisted as staff turned over. Without the log, each new scanner repeats the work of the previous one; with the log, the work accumulates and the scan's capability grows over time.

The log's accessibility determines whether it is used. The log must be accessible to the scanning team, to the foresight analysts who use it for foresight exercises, and to the strategy function and the planning authority for their own reference. The accessibility is managed: not every user has the same access, and the log's sensitive entries — those bearing on national security, market-moving information, or confidential interagency material — are protected by classification and access controls. The principle is that the log is a shared resource, not a private record, and that access controls are designed to enable use by those who need the log, not to prevent use by them. A log accessible only to the scanners who maintain it fails to function as institutional memory; the

mature practice maintains the log as a shared resource, with access controls appropriate to the content and with the documentation that allows users to find what they need.

8.4 The Review and Update Cycle

The review and update cycle is the periodic reassessment of the scan log and the scan's scope. The cycle ensures the log remains current, the assessments remain accurate, and the scan remains aligned with the planning function's needs as those needs evolve over the planning cycle. A log that is not reviewed accumulates stale entries and outdated assessments; a log that is reviewed remains a living record. The review is conducted on a defined schedule — monthly for active signals, quarterly for the full log, annually for the scan's scope and source list. The annual review of scope and source list is timed to feed into the planning cycle's foresight input, so that the scan's coverage is reassessed as each plan period begins.

The reassessment of signals in light of new information is the core of the review. Each active signal is reviewed: has new information arrived, and has the signal been confirmed, contradicted, or developed since the last assessment? The reassessment updates the signal's impact, uncertainty, and time horizon as appropriate, and records the new information in the assessment narrative so that the log preserves the history of the signal's development. Signals that have been confirmed may be escalated to the strategy function; signals that have been contradicted may be downgraded in priority; signals that have developed may be clustered into emerging issues for further assessment. The reassessment is the mechanism by which the log reflects the current state of the scan's knowledge rather than the state at the moment of the signal's first detection.

The archival of resolved signals is the mechanism by which the log remains manageable. A signal that has been resolved — the development has occurred, the implications have been addressed in the plan, or the signal has been conclusively contradicted — is archived. The archival removes the signal from the active log while preserving it in the record, so that the active log remains focused on signals that warrant continued attention. The archived signal remains searchable and can be reactivated if conditions change, and the archival entry records the basis for the resolution so that future scanners can understand why the signal was archived. The archival is the mechanism by which the active log remains focused on current signals while the full record is preserved for future reference and reactivation.

The reactivation of archived signals is the mechanism by which the log's memory is used. An archived signal that becomes relevant again — because conditions change, because a new question arises, or because a related signal appears — is reactivated and brought back into the active log with its previous assessment and history. The reactivation allows the scan to draw on its accumulated memory rather than starting afresh, and a signal archived years ago may be reactivated when a new development makes it relevant to the current planning question. Conducted consistently, the review and update cycle keeps the scan log functioning as a living institutional memory. A log that is reviewed remains current, accurate, and aligned with the planning function's needs; a log that is not reviewed accumulates stale entries and loses its value as memory.

8.5 Chapter Summary

The structured assessment of each signal applies a defined framework: category, potential impact, uncertainty, time horizon, and assessment narrative. The framework ensures every signal in the log is assessed to the same standard and that the log's entries are comparable across scanners and over time. The assessment narrative is the substantive content of the entry and the basis on which the signal will be assessed by the scan's users.

Prioritisation ranks assessed signals for further investigation or escalation. The prioritisation matrix plots signals on impact and uncertainty, supplemented by novelty and relevance criteria. The matrix determines the scanner's allocation of effort and is reviewed periodically to ensure it remains appropriate.

The scan log is the foresight function's institutional memory. Its structure — searchable, linked, organised — is what allows it to function as memory. The log preserves knowledge across personnel turnover, planning cycles, and changes of government. Its accessibility as a shared resource, with appropriate access controls, is what makes it usable in practice.

The review and update cycle is the periodic reassessment of the log and the scan's scope. The cycle reassesses active signals, archives resolved signals, and reactivates archived signals when they become relevant again. The cycle ensures the log remains current, accurate, and aligned with the planning function's needs. A log that is reviewed remains a living record; a log that is not reviewed accumulates stale entries and loses its value.

Chapter 9 — Scenario Theory and Types

Scenarios are structured narratives of plausible futures.

9.1 What Scenarios Are

A scenario is a structured narrative describing a plausible future, constructed to test the resilience of present decisions against the range of ways the future might develop. The scenario differs from a forecast by describing what could happen rather than projecting what will happen, and from a vision by describing a plausible future rather than advocating a desired one. A scenario set comprises a small number of scenarios — typically three to five — chosen to span the range of plausible futures rather than to identify the most probable one. The scenario's defining characteristic is the distinction between plausibility and probability: a scenario must be possible, given what is known, but it need not be likely. The distinction is the source of the scenario's value, because the futures that defeat a plan are often those that were possible but not judged likely.

The scenario takes narrative form because the future it describes must be vivid enough for the planner to assess its implications. The narrative sets out the conditions that obtain in the future, the events that have occurred, and the state of the world. The elements must be internally consistent, given the drivers and uncertainties that define the scenario. Internal consistency is the analytical core of the scenario: a future in which the conditions could not coexist, given the relationships among drivers, is broken and useless for testing. The threshold is plausibility: a scenario must be able to arise from present conditions, even under unusual resolutions of the uncertainties, but it need not be the most likely path.

The scenario's purpose is to allow the planner to ask, of each future in the set, whether the plan would succeed under that future and, if not, what the plan would need to change. A plan constructed against a single future, however well-chosen, is exposed to the conditions that future does not anticipate. The scenario set supplies the alternatives against which the plan is stress-tested. The output of the test is a plan that is explicitly robust across the range, or a plan with identified failure conditions and the responses that would be triggered if those conditions arose. The scenario is thus the instrument by which the plan's resilience is examined before conditions demand it, not after.

The distinction among scenarios, forecasts, and visions is the source of each one's value, and confusing them produces characteristic failures. A forecast used as a scenario collapses the range into a single line and conceals the conditions that would defeat the plan. A vision used as a scenario substitutes advocacy for analysis and produces a preferred future dressed as exploration. The Planning Manual treats scenarios at the general level in Chapter 7 and forecasting as an input to the baseline in Chapter 9; the doctrine develops the specialty craft of scenario theory, types, construction, and use. The Shell scenario tradition, developed continuously since the early 1970s, provides the most influential single body of comparative practice in the field and is the principal reference for the chapters that follow.

9.2 Exploratory, Normative, and Predictive Types

Exploratory scenarios start from the present and trace how the future might develop, without specifying a preferred endpoint. The construction identifies the drivers that will shape the future, assesses the critical uncertainties, and develops narratives that span the range of plausible developments. The exploratory scenario serves the planner’s principal question: what are the futures against which the plan must be robust. Exploratory scenarios dominate planning practice because the planner’s question is most often about resilience rather than about a target. The Shell global scenarios, the United Kingdom Government Office for Science Foresight projects, and the Netherlands Scientific Council for Government Policy’s prospectives are exemplars of exploratory practice at the national and sectoral scale.

Normative scenarios start from a desired future and work backwards to the present, identifying the conditions required for that future to arise and the path that would lead to it. The desired future is fixed by political direction, by public commitment, or by strategic objective. The normative scenario answers the planner’s question when the objective is given and the question is how to reach it: a plan to achieve universal primary education, a plan to reach carbon neutrality by a fixed date, a plan to eliminate a disease. Backcasting, the method associated with normative scenarios, was formalised in energy futures work in the 1970s and 1980s and is now standard in climate and sustainability planning. The normative scenario differs from the exploratory scenario in its direction of construction: backward from a defined future rather than forward from present conditions.

Predictive scenarios attach probabilities to a small number of futures, combining the exploration of possibilities with the assessment of likelihood. The type is the closest to forecasting among the three and is the most contested among scenario practitioners. Proponents value the explicitness about likelihood and the integration of exploration with probabilistic assessment; critics point to the weak basis for probability assignment at long horizons, the false precision that numerical probabilities can convey, and the risk that users collapse the set into its highest-probability member. The Intergovernmental Panel on Climate Change’s representative concentration pathways, used with probabilistic climate modelling, illustrate predictive practice at the frontier of the method’s defensible application. The type is used where the planning question requires likelihood as well as possibility and where the methods for probability assessment are defensible.

The choice of type is governed by the planning question, not by the practitioner’s preference or the institution’s habit. A national development plan, which must be robust against a range of futures, calls for exploratory scenarios. A sectoral plan with a fixed target — universal access to a service, elimination of a hazard, achievement of a standard — calls for normative scenarios to identify the path. A fiscal plan, where the uncertainty is largely economic and the economic models produce probability distributions, may call for predictive scenarios. The selection, made at the scoping stage, determines the methods, the participants, and the form of the output, and is treated at the general level in Chapter 7 of the Planning Manual and in specialty depth in the chapters that follow.

9.3 The Relationship to Forecasting and Visioning

Scenarios and forecasting address uncertainty through different operations. Forecasting quantifies uncertainty: the forecast produces a point estimate with a confidence interval, and the uncertainty is expressed as the range around the estimate. Scenarios explore uncertainty: the scenario set produces multiple futures, and the uncertainty is expressed as the differences among them. Forecasting is appropriate where uncertainty is well-characterised, the underlying processes are stable, and the range is narrow; scenarios are appropriate where uncertainty is deep, the processes are in flux, and the range is wide. Planning deals primarily with the latter kind of uncertainty, which is why scenarios are central to the discipline, but both are required: the forecast supplies the baseline, and the scenarios supply the resilience test.

Scenarios and visioning differ in their stance toward desirability. A vision describes a desired future and the path to it; a scenario describes a plausible future without judging its desirability. The vision is normative, expressing what the government aims to achieve, and the scenario is analytical, expressing what could arise. The two are complementary: the vision identifies the future the government works toward, while the scenarios identify the futures the government must be prepared to address, including those that would prevent the vision's realisation. A plan with a vision and no scenarios is unprepared for the conditions that could defeat it; a plan with scenarios and no vision lacks direction.

The complementarity of forecasting, scenarios, and visioning operates through the plan's structure. The forecast supplies the baseline: the most likely future against which the plan is constructed. The scenarios supply the resilience test: the range of futures against which the baseline is tested. The vision supplies the direction: the desired future toward which the plan aims. Together, the three constitute the plan's temporal framework — where conditions are heading, what conditions could arise instead, and what the plan is for — and the relationship is maintained through the planning cycle, with the forecast updated annually, the scenarios revisited at the start of each plan period, and the vision reviewed when political authority changes or when conditions demand it.

The integration is the planner's responsibility, and it depends on maintaining the distinction among the three. The forecast is produced by the forecasting function — the statistics office, the finance ministry, or the central bank; the scenarios are produced by the foresight function; the vision is set by political authority with planning support. The planner constructs the plan against the forecast, tests it against the scenarios, and aligns it with the vision. Each loses its value when collapsed into the others: a forecast used as a scenario conceals the range; a scenario used as a forecast produces false certainty; a vision used as a scenario substitutes advocacy for analysis. Planning authorities that maintain capacity in all three integrate them through the plan, and the discipline of maintaining the distinction is the condition under which the integration holds.

9.4 Scenario Sets

The scenario set is a small number of scenarios — typically three to five — chosen to span the range of plausible futures rather than to identify the most probable one. The set's logic rests on a judgment about how many futures a planner can hold in mind at once: a single scenario is a

forecast that collapses the range into a line, and two scenarios are a binary choice that misses the futures between or outside. Three to five scenarios capture the principal uncertainties without overwhelming the user; beyond five, the set loses clarity and the planner cannot work with all members at once. The set is the unit of scenario work, and the individual scenario is meaningful only as a member of a set. The set's design — the number of scenarios, the axes they span, and the members chosen for development — is the principal craft of the scenario phase and the subject of the assessment.

The logic of the set is governed by the critical uncertainties identified through driver mapping, treated in Chapter 10. The critical uncertainties are the drivers with high uncertainty and high impact — the forces whose future trajectory is unknown and whose effect on the system is large. These drivers define the axes along which the scenarios differ: two critical uncertainties produce a two-by-two matrix of four scenarios, and three produce a more complex structure of eight or more. Standard practice works with two critical uncertainties, producing four scenarios, of which three or four are developed, depending on which span the range most usefully. The choice of axes is the most consequential decision in scenario construction, because the axes determine what the set spans and what it misses.

Symmetric sets have scenarios that are equally developed and equally plausible: each scenario is given the same depth of narrative and the same weight in the assessment. The form is even-handed, requiring the planner to take each member seriously without privileging any. In practice, the form is demanding: in many exercises, one or two cells of the matrix are markedly less plausible than the others, and developing them to the same depth as the more plausible cells produces strain on the narrative. Symmetric sets suit planning questions where no future is more probable than the others and where the purpose is to test resilience across the full range. The Shell global scenarios and the Intergovernmental Panel on Climate Change's representative concentration pathways are exemplars of the symmetric form.

Asymmetric sets have a central scenario that is more developed and more probable, surrounded by variant scenarios that are less developed and less probable. The form reflects the realistic judgment that some futures are more probable than others, with the plan constructed against the most probable one and the variants providing the resilience test. The corresponding risk is that variant scenarios may be treated as outliers rather than as genuine possibilities, weakening the resilience test that is the set's purpose. Central-scenario-with-variants is the standard form in fiscal planning, in macroeconomic surveillance under the European Union's Stability and Growth Pact, and in much sectoral planning. The choice between symmetric and asymmetric sets is governed by the planning question and by the strength of the evidence for a central case.

9.5 Chapter Summary

Scenarios are structured narratives describing plausible futures, constructed to test the resilience of present decisions against the range of ways the future might develop. The scenario differs from the forecast — which projects what will happen — and from the vision — which describes a desired future. A scenario set is a small number of scenarios, typically three to five, chosen to span the range of plausible futures. The scenario's defining characteristic is the

distinction between plausibility and probability: the scenario must be possible, given what is known, but it need not be likely. The narrative must be internally consistent and vivid enough for the planner to assess the plan's implications under each future in the set.

Three principal types of scenarios are distinguished. Exploratory scenarios start from the present and explore how the future might develop; they dominate planning practice and address the question of plan resilience. Normative scenarios start from a desired future and work backwards; they are used in target-setting and in the assessment of strategic commitments. Predictive scenarios attach probabilities to a small number of futures; they are used where the planning question requires likelihood as well as possibility and where the methods for probability assessment are defensible. The choice of type is governed by the planning question, not by the practitioner's preference.

Scenarios, forecasting, and visioning are complementary. Forecasting addresses uncertainty by quantifying it; scenarios address it by exploring it; visioning addresses desirability by stating a desired future. The three together constitute the plan's temporal framework: the vision supplies direction, the forecast supplies the baseline, and the scenarios supply the resilience test. The integration is the planner's responsibility, and it depends on maintaining the distinction among the three, because each loses its value when collapsed into the others.

The scenario set is governed by the critical uncertainties identified through driver mapping, and the choice between symmetric and asymmetric sets depends on the planning question and the strength of the evidence for a central case. The set's construction is the foresight practitioner's principal craft: a small number of scenarios, typically three to five, chosen to span the range of plausible futures without overwhelming the user. The set's use in planning is to test the plan's resilience across the range, identifying the elements that are robust and the elements that require revision or contingent response. The craft of set construction is treated in operational detail in the chapters that follow.

Chapter 10 — Driving Force Identification and Analysis

Driving forces are the factors that will shape the future of the system under study.

10.1 The STEEPL Framework

The STEEPL framework is a six-category checklist for the identification of driving forces: social, technological, economic, environmental, political, and legal. The framework's function is to ensure that identification covers the full range of forces that could shape the future of the system under study, and to provide a structure against which completeness can be tested. Each category groups forces that operate through distinct mechanisms and that require distinct methods of assessment. The framework is treated at the general level in Chapter 8 of the Planning Manual; the doctrine develops the specialty craft of driver identification, assessment, and mapping.

Social forces comprise demographic change, social attitudes, and cultural change; they operate over long horizons and are among the most predictable, because demographic trajectories are established by present populations. Technological forces comprise research and development, technology adoption, and the diffusion of innovations; they operate over varying horizons and are among the most uncertain, because the pace and direction of change are difficult to project. Economic forces comprise growth, trade, finance, and the structure of the economy; they operate over medium horizons and are subject to cycles and shocks. The application of the framework to a particular question weights the categories by their relevance: an exercise on energy transition will emphasise technological, economic, and environmental forces; an exercise on social policy will emphasise social and political forces.

Environmental forces comprise climate change, resource depletion, ecological change, and natural hazards; they operate over long horizons and are among the most consequential, because their effects are large and difficult to reverse. Political forces comprise electoral trends, governance change, political instability, and geopolitical shifts; they operate over varying horizons and are among the most difficult to assess, because they depend on human decisions. Legal forces comprise regulatory change, judicial decisions, and international law; they operate over medium horizons and are closely linked to political forces. The assessor records each identified force in the category to which it best belongs, with the recognition that the classification is a working arrangement rather than an analytical conclusion.

The framework's limitations must be acknowledged. The categories overlap: a force may be social and economic, as in the effect of an ageing population on labour markets, or technological and political, as in the regulation of a new technology. The categories are not equally weighted in every exercise, and the framework provides no guidance on weighting; weighting is the assessor's judgment, informed by the planning question. The framework is a checklist, not a theory; it ensures completeness but does not explain the forces' interactions. The interactions are addressed through driver mapping, treated in section 10.4, and through cross-impact analysis, treated in Part IV.

10.2 Identification Methods

Driving forces are identified through three principal methods: literature review, expert elicitation, and structured workshops. Each method contributes distinct material, and the mature practice uses all three. A driver identification conducted by one method alone is vulnerable to that method's limitations: literature review alone misses forces not yet in the published record; expert elicitation alone misses forces outside the experts' fields; workshops alone miss forces not represented among the participants. The combination of methods is the condition under which the identification is comprehensive.

Literature review surveys the published record for forces identified by other analysts. The review covers the scientific literature, the technical press, government reports, and international organisation assessments; the International Energy Agency's World Energy Outlook and the OECD's International Futures Programme are examples of the latter. The review's strength is breadth: it draws on the work of many analysts across many fields. The review's weakness is lag: the forces in the published record are those that have been identified and reported, which may not include the most recent or the most novel. The review produces a list of forces with their sources, which enters the driver identification's working record.

Expert elicitation, treated in Chapter 6, brings experts' judgment to the identification. Experts in each STEEPL category are asked to identify the forces they consider most significant and to characterise each force's direction, strength, and time horizon. The elicitation's strength is depth: experts can identify forces that are not yet in the published record and can assess the forces' significance. The elicitation's weakness is the experts' biases: disciplinary biases, institutional biases, and personal biases shape what each expert identifies. The elicitation is managed through the methods treated in Chapter 6: diverse experts, structured elicitation, and documented dissent.

Structured workshops bring the foresight practitioners, the domain experts, and the planners together to identify forces collectively. The workshop's strength is integration: the participants identify forces, challenge each other's identifications, and develop a shared understanding of the forces' significance. The workshop's weakness is the risk of groupthink and the risk of dominance by the most assertive participants; both are managed by skilled facilitation, treated in Chapter 4. The output of the three methods is a list of driving forces, each with a description, a source, and a preliminary assessment. The list enters the driver map, treated in section 10.4.

10.3 Assessment of Direction, Strength, and Time Horizon

The assessment is the stage at which the list of identified forces becomes an analytical product. Each force is assessed for direction, strength, and time horizon, using a framework applied consistently across all forces. The consistency is the condition under which the driver map is a structured record rather than a collection of descriptions. The assessment is documented, so the map can be reviewed, challenged, and revised as new information arrives.

Direction is the tendency of the force: increasing, decreasing, or stable. A force that is increasing will have greater effect in the future; a force that is decreasing will have less; a force that is stable will continue at its present level. Direction is assessed on the basis of evidence: the trend in the data, the expert judgment, and the published assessments. Some forces have clear directions,

as population ageing in most mature jurisdictions; others are contested, as the direction of inequality or the direction of geopolitical instability. The assessment records the direction and the basis for it, including the dissent where the direction is contested.

Strength is the magnitude of the force's effect on the system under study. Strength is assessed on a defined scale — high, medium, low — with criteria for each level. A force with high strength will substantially shape the future; a force with low strength will have marginal effect. Strength is the assessor's judgment, informed by the evidence and by the force's relationship to the system: technological change is a strong force for an energy system, a weaker force for a social policy system. The assessment records the strength and the system to which it applies, so the same force may carry different strength ratings across exercises.

Time horizon is the period over which the force operates. Three bands are standard: near-term, within the plan horizon; medium-term, within ten years; long-term, beyond ten years. Some forces operate across all horizons, as demographic change; others operate at specific horizons, as a regulatory change that takes effect at a defined date. The time horizon is assessed on the basis of the force's characteristics and the evidence of its development. The assessment framework — direction, strength, time horizon — is applied to every force, and the assessment is documented to support review and revision.

10.4 The Driver Map

The driver map is the structured representation of the driving forces and their relationships. The map comprises the list of forces, each with its assessment, and the representation of the relationships among them. The relationships are the connections among forces: forces that reinforce each other, forces that counteract each other, and forces that interact to produce effects neither would produce alone. The map's purpose is to provide the analytical foundation for scenario construction.

The map's structure has two layers. The first layer is the list of forces, organised by STEEPL category, with each force's assessment recorded. The second layer is the representation of relationships: which forces affect which, in what direction, and with what strength. The relationships are identified through cross-impact analysis, treated in Part IV, and through expert judgment. The representation may be a matrix, with forces on both axes and the cells indicating the relationship; a diagram, with forces as nodes and relationships as arrows; or a narrative, with the relationships described in text, the choice among them governed by the map's audience and purpose.

The map's principal use is to support scenario construction. The scenarios are built from the critical uncertainties — the forces with high uncertainty and high impact, identified from the map — which define the axes along which the scenarios differ, treated in Chapter 9. The map also identifies the predetermined elements: the forces whose direction is clear and whose development is largely certain. The predetermined elements are common across all scenarios; the critical uncertainties differ among them. The map's quality determines the scenarios' quality: a map with weak identification or weak assessment produces scenarios that are not grounded in the analysis.

The map's use extends beyond scenario construction. The map informs the scan's scope, by identifying the forces the scan should monitor; the risk register, by identifying the forces that could produce risks; and the plan's assumptions, by identifying the forces the plan takes as given. The map is the foresight function's analytical product, and its maintenance is the function's ongoing responsibility. The map is reviewed and updated as new information arrives, as forces develop, and as the system under study changes; a map that is not updated loses its value, while a map that is maintained remains the analytical foundation for the foresight function's work.

10.5 Worked Example: Driver Identification for Energy Transition

Consider a foresight exercise on the energy transition over a twenty-year horizon, conducted by a foresight unit in support of a national energy plan. The exercise's question is: what conditions could shape the energy system over the next twenty years, and what are the implications for the plan? The foresight unit convenes a team comprising a foresight analyst, two domain experts (one in energy technology, one in energy economics), and a facilitator. The team is supported by the scanning team's scan log, which provides the accumulated signals on energy developments.

The team begins with literature review. The review covers the scientific literature on energy technology, the technical press on energy markets, the international organisation assessments (the International Energy Agency's World Energy Outlook, the International Renewable Energy Agency's reports), and government reports from jurisdictions pursuing energy transitions. The review identifies a list of forces: the declining cost of renewable generation, the development of battery storage, the electrification of transport, the regulatory pressure on fossil fuels, the geopolitical shifts in energy supply, the investment requirements for grid modernisation, and the social acceptance of energy infrastructure. Each force is recorded with its source.

The team conducts expert elicitation with a panel of ten experts drawn from academia, industry, and government. The experts are asked to identify the forces they consider most significant and to assess each force's direction, strength, and time horizon. The elicitation identifies additional forces: the development of hydrogen as an energy carrier, the role of nuclear power, the demand-side effects of efficiency improvements, and the geopolitical implications of energy independence. The elicitation also surfaces disagreement: the experts disagree on the pace of renewable cost reduction, on the role of nuclear, and on the demand-side effects. The disagreement is documented and enters the driver map as uncertainty.

The team convenes a structured workshop with the experts, the planners, and the stakeholders. The workshop reviews the list of forces, challenges the assessments, and identifies the relationships among the forces. The workshop identifies the critical uncertainties: the pace of technological development, the strength of regulatory pressure, and the social acceptance of infrastructure deployment. The workshop also identifies the predetermined elements: the demographic growth that will drive demand, the existing infrastructure that will shape the transition path, and the international commitments that constrain the options.

The team produces the driver map. The map lists the forces by STEEPL category, with each force's assessment. The map represents the relationships: the declining cost of renewables reinforces the electrification of transport; the regulatory pressure reinforces the investment in grid

modernisation; the social acceptance constrains the deployment of infrastructure. The map identifies the critical uncertainties and the predetermined elements. The map enters the scenario construction phase, in which the team builds a scenario set spanning the range of plausible energy futures, treated in Chapter 9.

10.6 Chapter Summary

The STEEPL framework is the standard structure for identifying driving forces. The framework ensures that the identification covers the full range of forces and provides a checklist for completeness. The categories overlap, are not equally weighted in every exercise, and constitute a checklist rather than a theory. The interactions among forces are addressed through driver mapping and cross-impact analysis.

Driving forces are identified through three methods: literature review, expert elicitation, and structured workshops. Each method contributes distinct material, and the mature practice uses all three. Literature review provides breadth; expert elicitation provides depth; workshops provide integration. The combination is the condition under which the identification is comprehensive.

Each force is assessed for direction, strength, and time horizon. The assessment framework is applied consistently across all forces, and the assessment is documented with the evidence and the reasoning that support it. The documentation allows the driver map to be reviewed, challenged, and revised as new information arrives. The discipline of recording the basis for each assessment is what distinguishes a driver map from an opinion piece, and what permits the map to be maintained across staff turnover and exercise cycles.

The driver map is the structured representation of the forces and their relationships. The map’s purpose is to support scenario construction: the scenarios are built from the critical uncertainties and the predetermined elements identified from the map. The map also informs the scan’s scope, the risk register, and the plan’s assumptions. The worked example on energy transition illustrates the full process from identification through assessment to the production of the driver map.

Chapter 11 — Critical Uncertainty Analysis

Part III: Scenario Methods — The construction of alternative futures against which strategy is tested.

11.1 What Critical Uncertainty Analysis Is

Critical uncertainty analysis is the structured assessment of the driving forces that shape the future of the system under study, conducted to identify the small number of forces whose future trajectory is most uncertain and whose effect on the system is largest. A driving force is any factor that influences the development of the system; uncertainty is the degree to which the force's future trajectory is unknown; impact is the magnitude of the force's effect on the system. The analysis combines the two assessments to produce a two-dimensional classification that separates forces which are predictable and marginal from those which are unpredictable and consequential. The forces in the latter category are the critical uncertainties, and they become the axes along which the scenarios differ.

The assessment is conducted over the full list of driving forces identified through horizon scanning, expert elicitation, and the methods treated in Chapter 10. Each force is assessed for uncertainty and for impact, on scales applied consistently across the set, with the basis for each rating recorded. The assessment draws on the evidence assembled in the driver map and on the judgment of the practitioners, the domain specialists, and the stakeholders engaged in the exercise. The output is a classified register, in which each force occupies a position in the uncertainty-impact space and the critical uncertainties are visible as the cluster in the high-uncertainty, high-impact quadrant.

The critical uncertainties are those forces that combine high uncertainty with high impact. A force with high impact but low uncertainty is a predetermined element: its effect on the system is large, but its trajectory is sufficiently clear that it need not vary across scenarios. A force with high uncertainty but low impact is a secondary driver: its trajectory is unknown, but its effect is too small to anchor the differences among scenarios. Only the forces that are high on both dimensions qualify as critical, because only these produce the variation the planner must hold in mind at once when constructing and testing the set.

The relationship to scenario construction is direct. The critical uncertainties define the axes of the scenario matrix, and the combinations of their resolutions define the cells. Two critical uncertainties produce a two-by-two matrix of four scenarios; three produce a more complex structure of eight or more cells. The number chosen governs the structure of the set, the depth of the narratives, and the range the set spans. The critical uncertainty analysis is therefore the bridge between the driver mapping treated in Chapter 10 and the scenario construction treated in Chapter 9: it selects, from the full inventory of forces, the subset that the scenario exercise will work with.

11.2 The Uncertainty-Impact Assessment

The assessment is two-dimensional. The first dimension is uncertainty: the degree to which the force's future trajectory is unknown, assessed at the planning horizon. A force is highly uncertain when the evidence does not constrain its trajectory, when expert judgment diverges, or when the force depends on decisions that have not yet been taken. A force is low in uncertainty when its trajectory is set by structural conditions, by committed decisions, or by well-established trends. Population ageing in most mature jurisdictions is low-uncertainty, because the demographic trajectory is established by present populations; the pace of a specific technological breakthrough is high-uncertainty, because the timing depends on research outcomes not yet achieved.

The second dimension is impact: the magnitude of the force's effect on the system under study. Impact is assessed relative to the system, not in absolute terms: a force that is consequential for an energy system may be marginal for a health system, and a force that is consequential for a national plan may be marginal for a sectoral plan. The assessment is informed by the driver map's representation of relationships, by the magnitude of the effect documented in the literature, and by the judgment of those who work with the system. Impact is rated on a defined scale — high, medium, and low — with criteria for each level applied consistently across all forces.

The assessment method is structured to produce consistent, comparable judgments across the full list of forces. Each force is assessed by the foresight team in consultation with the domain experts, using the evidence assembled in the driver map and any additional evidence the exercise has gathered. The assessment is recorded on a scoring sheet or in a structured register, with the uncertainty rating, the impact rating, and a brief statement of the basis for each. Recording the basis is the discipline that separates the assessment from an opinion: it makes the judgment reviewable, allows disagreement to be documented, and permits the assessment to be revised as new evidence arrives.

The documentation of the assessment comprises the full register, the supporting evidence for each rating, and the recorded dissent where experts disagreed. The register is maintained as part of the driver map, with cross-references between the uncertainty-impact ratings and the relationships among forces recorded in the map. The documentation is the product the analysis leaves behind: it is the record on which the scenario construction rests and the record to which the foresight function returns when the exercise is revisited, reviewed, or revised. The register's maintenance across cycles is what permits an institution to build accumulated analysis rather than to begin each exercise from a blank sheet.

11.3 Identifying the Critical Uncertainties

The identification of the critical uncertainties is the selection of the forces that are high on both dimensions, drawn from the classified register produced by the uncertainty-impact assessment. The selection requires judgment: a force may sit near the threshold on one dimension or the other, and the decision of whether it qualifies as critical rests with the practitioner, informed by the evidence and the planning question. The selection is documented with the same discipline as the assessment, with the basis for each critical-uncertainty designation and each exclusion

recorded. The selected forces become the axes of the scenario matrix, and the selection's quality governs the set's quality.

The number of critical uncertainties is typically two. Two axes produce a two-by-two matrix of four scenarios, of which three or four are developed, depending on which span the range most usefully for the planning question. The number is a working convention, not a theoretical constraint: it reflects the judgment that a planner can hold two axes of variation in mind at once, that four scenarios span a useful range without overwhelming the user, and that the matrix structure supports clear narrative development. The convention is observed across the principal exemplars of practice, including the Shell global scenarios, the United Kingdom Government Office for Science Foresight projects, and the Netherlands Scientific Council for Government Policy's prospectives.

The trade-off between complexity and clarity governs the choice of number. Three critical uncertainties produce a matrix of eight cells, which captures more of the variation in the system but strains the planner's ability to work with the full set. One critical uncertainty produces a binary contrast, which is clear but misses the futures that lie between or outside the two extremes. The choice is governed by the planning question and by the structure of the uncertainty in the system: a question that turns on two principal uncertainties is well-served by a two-by-two matrix, while a question that turns on three may justify the additional complexity. The selection is revisited at the start of each exercise, and the choice made for one exercise does not bind the next.

11.4 The Two-by-Two Matrix and Alternatives

The two-by-two matrix is the standard configuration produced by two critical uncertainties. Each axis represents the resolution of one critical uncertainty, with the endpoints defined by the two principal ways the uncertainty could resolve. The four cells of the matrix are the scenario seeds: each cell is a combination of resolutions, and each combination defines a future that the scenario narrative will develop. The matrix's principal property is symmetry: the four cells are constructed by the same operation, and each is treated as a serious possibility, with no cell privileged as the central case unless the exercise is explicitly designed as central-with-variants.

The four cells, developed into narratives, become the scenario set. The narratives must be internally consistent: each cell's combination of resolutions must be plausible, given the predetermined elements and the relationships among the forces. The narratives must also be distinct: each cell must produce a future that is recognisably different from the others, so that the resilience test the set provides is genuine. Where two cells produce futures that are too similar to be distinguished, the exercise has identified a redundancy in the matrix, and the practitioner must either revise the axes or merge the redundant cells. The Shell global scenarios and the Intergovernmental Panel on Climate Change's representative concentration pathways are exemplars of the four-cell matrix developed into a working scenario set.

Alternatives to the two-by-two matrix are used when the structure of the uncertainty in the system does not fit the symmetric form. A three-by-three matrix, with three levels of each axis, produces nine cells and captures intermediate outcomes that the binary axes miss. The configuration is used when the critical uncertainties have a meaningful middle state — a moderate pace of technological change, a partial regulatory response — that the planner needs to explore.

The cost is the loss of clarity that comes with nine scenarios: the set is harder to communicate, harder to work with, and harder to keep internally consistent across the cells.

Asymmetric matrices and non-matrix approaches are used in further departures from the standard form. An asymmetric matrix has axes of unequal resolution, as when one uncertainty is binary and another has three levels, producing six cells; the form is used when the two critical uncertainties have different natural structures. A non-matrix approach identifies a central scenario and a small number of variants, each varying one critical uncertainty from the central case; the form is used when the central case is well-supported and the planning question concerns the resilience of that case rather than the exploration of the full range. Each configuration has its appropriate use, and the choice among them is governed by the planning question, the structure of the uncertainty, and the audience for the scenarios.

11.5 The Relationship to the Driver Map

The critical uncertainty analysis is a product of the driver mapping treated in Chapter 10. The driver map supplies the inventory of forces, the assessment of direction and strength, and the representation of relationships among forces. The uncertainty-impact assessment builds on the map: the impact rating draws on the strength assessment recorded in the map, and the uncertainty rating draws on the divergence among experts and the evidence assembled in the map's documentation. An analysis conducted without the map rests on an undocumented list; an analysis conducted with the map rests on the accumulated judgment of the foresight function.

The driver map informs the uncertainty-impact assessment in three principal ways. The strength assessment recorded for each force in the map provides the starting point for the impact rating: a force rated high in strength in the map is a candidate for high impact in the assessment, with the rating confirmed or revised on the basis of the assessment's specific scope. The relationships recorded in the map identify the forces that act through other forces, which may amplify or attenuate the impact a force would have in isolation. The cross-impact analysis, where it has been conducted, identifies the forces whose resolution would propagate most widely through the system, which is itself a measure of impact.

The documentation of the linkage between the driver map and the critical uncertainty analysis is what allows the analysis to be reviewed and revised. Each critical uncertainty's designation is cross-referenced to the force's record in the driver map, including the strength assessment, the relationships, and the expert judgments. Each exclusion is cross-referenced in the same way, with the basis for the exclusion recorded. The linkage makes the analysis traceable: a reviewer can follow a critical uncertainty from its designation back to the underlying force, the assessment, and the evidence, and the traceability is what separates the analysis from a selection made on intuition. The traceability is also the property that allows the analysis to be maintained across staff turnover and exercise cycles, since a successor can read the record and resume the work.

11.6 Chapter Summary

Critical uncertainty analysis is the structured assessment of driving forces to identify the small number whose future trajectory is most uncertain and whose effect on the system is largest. The analysis combines the assessment of uncertainty with the assessment of impact, producing a two-dimensional classification that separates forces which are predictable and marginal from those which are unpredictable and consequential. The forces in the latter category are the critical uncertainties, and they define the axes along which the scenarios differ. The analysis is the bridge between driver mapping and scenario construction: it selects, from the full inventory of forces, the subset that the scenario exercise will work with.

The uncertainty-impact assessment is two-dimensional. Uncertainty is the degree to which the force's future trajectory is unknown; impact is the magnitude of the force's effect on the system, assessed relative to the system under study. Each force is assessed on both dimensions, on scales applied consistently across the set, with the basis for each rating recorded in a structured register. Recording the basis is the discipline that separates the assessment from an opinion: it makes the judgment reviewable, allows disagreement to be documented, and permits the assessment to be revised as new evidence arrives.

The critical uncertainties are the forces high on both dimensions. The number is typically two, producing a two-by-two matrix of four scenarios; the convention reflects the judgment that a planner can hold two axes of variation in mind at once, that four scenarios span a useful range without overwhelming the user, and that the matrix structure supports clear narrative development. The trade-off between complexity and clarity governs the choice: three axes capture more variation but strain the planner's capacity; one axis is clear but misses the futures between or outside the extremes. The choice is governed by the planning question and the structure of the uncertainty in the system, and it is revisited at the start of each exercise.

The two-by-two matrix is the standard configuration, with each axis representing the resolution of one critical uncertainty and each cell a scenario seed. Alternatives — three-by-three matrices, asymmetric matrices, and non-matrix approaches — are used when the structure of the uncertainty does not fit the symmetric form, with each configuration matched to a different planning question and audience. The critical uncertainty analysis is a product of the driver map, which supplies the inventory of forces, the strength assessment, and the relationships among forces that the analysis builds on. The documentation of the linkage between the analysis and the map is what allows the analysis to be reviewed, revised, and maintained across exercise cycles, and it is the property that distinguishes institutional analysis from individual judgment.

Chapter 12 — Scenario Construction

The development of scenario narratives from the critical uncertainty matrix.

12.1 What Scenario Construction Is

Scenario construction is the discipline by which the foresight analyst develops narrative futures from the critical uncertainty matrix. Where the matrix identifies the axes along which the future may differ, the construction translates the cells of the matrix into coherent descriptions of futures that could plausibly emerge. The output of construction is a set of scenarios, rather than a single projected future; each scenario is internally consistent, and each differs from the others on the critical uncertainties that define the matrix. The discipline extends the general doctrine of scenario methods set out in Chapter 7 of the Planning Manual; the doctrine develops the craft of construction itself.

The translation of a matrix cell into a scenario narrative is the central act of construction. Each cell defines a combination of outcomes for the critical uncertainties; the analyst's task is to specify the future that would result from that combination, given the predetermined elements and the relationships among the drivers recorded in the driver map. The cell defines the structure; the narrative develops the content, supplying the events, the conditions, and the implications that distinguish this future from the others in the set. The narrative must account for the drivers the map has identified, must respect the relationships the cross-impact analysis has established, and must remain consistent with the predetermined elements that are common across all scenarios.

Scenario construction is distinct from forecasting and from creative writing, in ways treated in Chapter 9. A forecast selects the most probable future and projects it; scenario construction develops a set of futures without ranking them by probability. A work of fiction invents its world from the author's imagination; a scenario is bound to the evidence on drivers, to the relationships in the driver map, and to the logic that connects them. Pierre Wack, who led the scenario practice at Royal Dutch Shell in the 1970s and 1980s, framed the working distinction: scenarios are constructed to change the decision-maker's mindset, while forecasts are constructed to predict the future. The construction serves that purpose by producing futures that are plausible, vivid, and challenging to the assumptions embedded in present decisions.

12.2 The Scenario Narrative

The scenario narrative comprises five components: the title, the opening, the trajectory, the end-state, and the implications. The title condenses the scenario's character into a short phrase; the opening sets the scene at the horizon date and locates the reader in the future; the trajectory traces the path from the present to that future, identifying the events and decisions that mark the path; the end-state describes the conditions that obtain at the horizon date; the implications state what the future means for the planning question. The five components are standard across mature practice and are visible in the Shell scenario publications and in the Global Business Network's template for

scenario reporting. Each component performs a distinct function, and the omission of any weakens the scenario's capacity to inform decisions.

The narrative arc connects the components and gives the scenario its coherence. The arc runs from the present through a sequence of developments to the end-state, and the trajectory is the arc's load-bearing element. The trajectory is a plausible chain of events, identifying the kinds of events that would have to occur, the order in which they would occur, and the conditions under which they would occur. The arc must be neither too smooth nor too dramatic; a trajectory without turning points reads as an extrapolation, while a trajectory of constant shocks strains credibility. The Shell scenarios of the 1980s, which traced the oil price collapse through the interplay of supply growth and demand destruction, are a studied example of an arc with the right number of turning points in the right places.

The level of detail in the narrative is governed by the scenario's purpose. A scenario constructed to test a national plan must develop the conditions that bear on the plan's commitments in sufficient detail that the test is meaningful; it may treat other conditions summarily. A scenario constructed to inform a sectoral strategy must develop the conditions specific to the sector, with adjacent sectors sketched only where they impinge. The detail is governed by the decision the scenario is built to inform, and the analyst calibrates the level of detail to that decision. Excessive detail obscures the strategic content; insufficient detail leaves the test unperformable.

The scenario narrative avoids prescription. It describes a future and the path to it; it does not recommend the future, endorse the path, or judge the desirability of the end-state. A scenario that prescribes operates as a vision, and the analytical purpose of the scenario is forfeited. The discipline of avoiding prescription is hard, because the analyst often holds views on which future is preferable, and those views can shape the narrative subtly. The mature practice manages this by separating the construction team from the evaluation team, by reviewing scenarios against the brief for prescriptive language, and by testing the set as a whole for symmetry of desirability across the members.

12.3 Internal Consistency

Internal consistency is the requirement that the elements of the scenario cohere. Each claim the narrative makes about the future must be compatible with the other claims it makes, with the predetermined elements, and with the relationships among drivers recorded in the driver map. A scenario in which the demand for a service rises while the population that consumes it falls, without an intervening mechanism that explains the divergence, is internally inconsistent. The consistency is the condition under which the scenario is a plausible future rather than a list of independently asserted claims. Consistency is tested, not assumed; the testing is the second act of construction, following the initial drafting of the narrative.

The testing of consistency is conducted against the driver map. Each driver in the map has an assessment of direction and strength, and the scenario's treatment of each driver must be compatible with that assessment or must explicitly depart from it with stated reasons. The predetermined elements, which are common across all scenarios, must appear in every scenario in the same form. The relationships among drivers, identified through cross-impact analysis, must be

respected: where two drivers reinforce each other, the scenario that increases one must increase the other, or must explain why the reinforcement does not operate in this case. The testing produces a record of consistencies and inconsistencies, and the inconsistencies are resolved before the scenario is finalised.

Inconsistencies are resolved in three ways. The first is revision: the analyst revises the narrative so that the inconsistent element is removed or modified. The second is elaboration: the analyst adds an intervening mechanism that reconciles the elements, drawing on a driver relationship not previously specified. The third is exclusion: the analyst removes an element from the scenario, acknowledging that the combination of outcomes in that cell cannot be made consistent, and that the scenario fails the plausibility test on its own terms. The third outcome is the most consequential, because it requires the analyst to revisit the critical uncertainty matrix itself, with cross-impact analysis, treated in Chapter 15, providing the structured representation of driver relationships against which the testing and the resolution are conducted.

12.4 Vividness and Accessibility

The scenario must be vivid enough to inform decisions and accessible enough to be understood by the audience for which it is constructed. Vividness is the quality by which the scenario concretises the future in the reader’s mind, so that the implications for present decisions become apparent. A scenario that is abstract, that states outcomes without specifying their texture, leaves the decision-maker with nothing to react to. Vividness is achieved through detail, through example, and through contrast: detail that specifies the conditions of the future, examples that illustrate what those conditions would mean in practice, and contrast that sets the scenario against the present and against the other scenarios in the set. The Shell scenarios achieved vividness through the device of a future historian looking back from the horizon date, which allowed the conditions of the future to be described in concrete and particular terms.

The scenario must be accessible to non-specialists. The audience for a scenario set includes the political authority that owns the plan, the line ministries that deliver it, and the stakeholders affected by it; few of these readers are foresight specialists. The narrative must therefore be written in plain language, must define its technical terms on first use, and must avoid the jargon of foresight methodology. The Global Business Network’s published scenarios, which were written for senior executives without foresight training, are a useful reference for the register: the analytical apparatus is kept to the appendices, and the narratives themselves are written as if for an intelligent general reader. Accessibility is the condition of the scenario’s effectiveness, because a scenario that is not understood cannot inform a decision.

The balance between detail and clarity is the analyst’s craft. Too much detail produces a narrative that is hard to read and that obscures the strategic content; too little produces a narrative that is clear but uninformed. The balance is struck by the principle of relevance: detail is included where it bears on the planning question, omitted where it does not. The principle is easy to state and hard to apply, because the question of what bears on the planning question is itself a judgment. The principle is applied through revision: the first draft includes everything the analyst considers

relevant, the second draft cuts what subsequent reflection shows to be peripheral, and the third draft polishes the prose so that what remains is both detailed and clear.

12.5 The Scenario Set

The scenario set is the analytical product. A single scenario is a description of a future; the set is a representation of the range of plausible futures. The set's quality is judged by its coverage of the range, by the symmetry or asymmetry of its members, and by the number of scenarios it contains. The individual scenarios are constructed to serve the set, and the construction of any one is constrained by the need to maintain the set's properties. The set is the unit that enters the planning process, that supports the testing of the plan, and that is maintained over time.

Coverage of the plausible range is the set's primary property. The set must span the range of futures that could plausibly emerge from the critical uncertainties, without redundancy among the members and without gaps that exclude a plausible combination. The symmetry or asymmetry of the set is the second property: a symmetric set has members that are equidistant from the centre of the range, while an asymmetric set has a central case with variants on either side. The number of scenarios is the third property: three to five is the standard range, with four being the modal case produced by a two-by-two matrix. Numbers above five strain the capacity of decision-makers to engage with each scenario; numbers below three risk representing the range inadequately.

The set is managed over time. Scenarios are not constructed once and then held; they are reviewed against new evidence, against the development of the drivers, and against the trajectory of the world. A scenario that is being realised is identified through the monitoring system, and the implications for the plan are drawn out. A scenario that has become implausible is retired, with the reasoning documented, and a scenario that has emerged as plausible but was not in the original set is added, with the reasoning documented and the implications for the set's coverage assessed. The Shell practice, which has maintained a scenario set for over fifty years, treats the set as a living product of the foresight function, with revisions on a regular cycle and ad-hoc revisions in response to major events.

12.6 Worked Example: Constructing a Scenario Set

Consider a foresight exercise conducted by a national planning authority on the future of the workforce over a fifteen-year horizon, in support of a national skills strategy. The exercise has identified two critical uncertainties through the methods treated in Chapter 11: the pace of automation, which ranges from rapid to slow; and the responsiveness of the education and training system, which ranges from high to low. The two-by-two matrix produces four cells, each to be developed into a scenario. The predetermined elements include the demographic trajectory, the existing structure of the workforce, and the international commitments on labour mobility. The driver map, produced through the methods of Chapter 10, identifies the relationships among the drivers and is the reference against which the scenarios will be tested for consistency.

The first cell combines rapid automation with high responsiveness. The scenario, titled "Reskilled Transition", opens at the horizon date with a workforce in which the share of

employment in routine tasks has fallen by a third, while the share in non-routine cognitive tasks has risen commensurately. The trajectory traces the path: the early years see the deployment of automation technologies in routine-intensive sectors, the education and training system responds with new curricula and expanded capacity, and the workforce transitions through a combination of initial education, mid-career retraining, and occupational mobility. The end-state describes a workforce that has absorbed the shock of automation with moderate disruption. The implications for the skills strategy concern the level of investment in retraining capacity and the speed of curricular reform.

The second cell combines rapid automation with low responsiveness. The scenario, titled “Disrupted Transition”, opens at the horizon date with substantial displacement, persistent skills shortages in growing occupations, and elevated long-term unemployment among workers displaced from routine tasks. The trajectory traces the path: automation proceeds at the same pace as in the first scenario, but the education and training system does not expand capacity or reform curricula at the required speed. The end-state describes a workforce bifurcated between high-skill workers in growing occupations and displaced workers in precarious or absent employment. The implications for the skills strategy concern the treatment of displaced workers, the role of social protection, and the case for more directive intervention in the training system.

The third cell combines slow automation with high responsiveness, and the fourth combines slow automation with low responsiveness. The third scenario, titled “Prepared Stability”, describes a future in which the anticipated disruption does not materialise on the projected schedule, the investments in retraining capacity are in part surplus to requirement, and the workforce continues to absorb routine-task employment at near-present levels. The fourth scenario, titled “Drift”, describes a future in which neither automation nor training reform proceeds at pace, and the workforce continues on its present trajectory with the structural problems of the present unaddressed. The four scenarios are tested for internal consistency against the driver map, are reviewed for prescriptive language, and are calibrated for level of detail. The set enters the skills strategy as the framework against which the strategy’s commitments are tested for robustness, with the use of the set in strategy treated in Chapter 13.

12.7 Chapter Summary

Scenario construction is the discipline by which narrative futures are developed from the critical uncertainty matrix. The discipline translates the cells of the matrix into coherent scenarios, each internally consistent and each differing from the others on the critical uncertainties. Construction is distinct from forecasting in its handling of uncertainty and from creative writing in its binding to evidence. The scenario narrative comprises the title, the opening, the trajectory, the end-state, and the implications, connected by the narrative arc and held to a level of detail calibrated to the planning question.

Internal consistency is the condition under which a scenario qualifies as a plausible future. Consistency is tested against the driver map and against the relationships among drivers identified through cross-impact analysis. Inconsistencies are resolved by revision, by elaboration, or by exclusion, with the last requiring the analyst to revisit the critical uncertainty matrix itself.

Vividness and accessibility are the conditions under which a scenario can inform decisions: vividness is achieved through detail, example, and contrast; accessibility is achieved through plain language and the relegation of methodological apparatus.

The scenario set is the analytical product, and its quality is judged by coverage of the plausible range, by symmetry or asymmetry, and by number. The standard set contains three to five scenarios, with four being the modal case. The set is managed over time, with scenarios reviewed, retired, and added in response to new evidence and to the development of the world. The worked example on the future of the workforce illustrates the construction of a four-scenario set from two critical uncertainties, with the narrative development for each cell and the testing for consistency against the driver map.

Chapter 13 — Scenario Use in Strategy and Planning

Scenarios are constructed to be used.

13.1 What Scenario Use Is

Scenario use is the application of constructed scenario sets to the decisions the planner must make. The decisions include the formulation of strategy, the testing of plans, the selection of assumptions, and the design of contingency responses. A scenario set that is constructed but not applied is an analytical exercise without consequence; the value of the scenario method is realised only when the set informs a decision. Use is therefore the stage at which the investment in construction is converted into planning value. The planner's responsibility at this stage is to translate the analytical product of the foresight function into commitments that the planning authority can adopt, fund, and deliver.

Scenario construction, treated in the preceding chapter, produces the scenario set: the internally consistent narratives, the critical uncertainties, and the predetermined elements. Scenario use takes the constructed set as its input and applies it to the planner's questions. The two activities are sequential and dependent: use without construction is impossible, and construction without use is incomplete. the doctrine treats them in separate chapters because the craft differs: construction is analytical and imaginative, while use is deliberative and decision-oriented. A practitioner who is skilled in construction is not thereby skilled in use; the two competencies are developed separately and staffed accordingly.

The sequence carries logical force beyond procedural preference. The scenario set must be settled before it is applied: if the set changes during use, the use is invalidated. The discipline requires that the construction record — the driving forces, the critical uncertainties, the narrative logic — be available to those who use the scenarios, so that the use is grounded in the construction's reasoning. Where use is attempted against an undocumented or unstable set, the result is contention rather than decision. The construction record is therefore transmitted alongside the scenario set, and the planning authority that receives the scenarios receives the reasoning that produced them.

The discipline of scenario use is recognised across mature practice. Shell's scenario team applies its sets to strategy reviews; the RAND Corporation applies scenario exercises to defence and resource decisions; the OECD's International Futures Programme applies scenarios to policy assessments. Singapore's Centre for Strategic Futures and the United Kingdom Government Office for Science institutionalise the sequence by separating construction exercises from planning applications. The European Commission's Better Regulation Toolbox treats scenario application as a distinct step in impact assessment, subsequent to scenario development. The handbook extends this common practice rather than introducing a novel arrangement.

13.2 Testing Plan Resilience

Testing plan resilience is the application of the scenario set to the plan to determine whether the plan would succeed under each plausible future. The test is conducted scenario by scenario: the planner asks, for each scenario, whether the plan's objectives would be achieved, its instruments would function, and its commitments would hold. A plan that succeeds under all scenarios is robust; a plan that succeeds under some and fails under others is partially robust; a plan that fails under most has not been properly stress-tested. The output is a resilience profile of the plan against the scenario set. The profile is recorded in a structured form, so that the test can be repeated when the plan or the scenario set is revised.

The resilience test distinguishes two categories of plan elements. Robust elements are those that succeed under every scenario: the predetermined elements in the scenario set support these, because the conditions they describe hold across the range. Scenario-dependent elements are those that succeed under some scenarios and fail under others: these correspond to the critical uncertainties, where the conditions differ across the range. The planner records, for each plan element, the scenarios under which it succeeds and the scenarios under which it fails, and the record is the analytical basis for plan modification. The classification is revisited when the scenario set is revised, because an element that was robust under one set may be scenario-dependent under another.

Modification takes three forms. First, the planner may redesign scenario-dependent elements to make them robust, typically by reducing the element's dependence on the critical uncertainty. Second, the planner may convert fixed commitments into conditional commitments, with the condition tied to indicators of which future is emerging. Third, the planner may retain the element as scenario-dependent, document the conditions under which it would fail, and refer the failure conditions to the contingency-planning process. The choice among the three is a judgment informed by the cost of robustness, the feasibility of conditionality, and the political acceptability of failure, and the choice is recorded so that the reasoning can be reviewed when the scenario set is revised.

Plan-resilience testing is institutionalised in several mature jurisdictions. The Netherlands Environmental Assessment Agency subjects long-term plans to scenario stress tests; the European Commission's Recovery and Resilience Facility requires member-state plans to be tested against adverse scenarios; the United Kingdom's National Infrastructure Commission publishes resilience assessments of infrastructure plans. The common practice is to publish the resilience profile alongside the plan, so that the political authority and the public can see where the plan is robust and where it depends on conditions. the doctrine treats this transparency as a requirement of the discipline, not as a discretion of the planner. A plan whose resilience profile is withheld has not been fully tested, because the test's conclusions have not been exposed to challenge.

13.3 Identifying Indicators

Scenario indicators are the observable measures that distinguish which of the constructed scenarios is emerging. Each scenario describes a future that differs from the others on the critical uncertainties; the indicators are the measurable proxies for those uncertainties. The indicators are

identified by working backward from each scenario's defining conditions to the variables whose movement would signal those conditions. The output is a set of indicators, each linked to one or more scenarios, that the monitoring system can track. An indicator that does not discriminate among scenarios is of no use for this purpose, however interesting it may be for other monitoring ends.

Leading indicators change before the scenario's defining conditions materialise; they provide early signal that a future is emerging. Lagging indicators change after the conditions materialise; they confirm that a future has emerged. The distinction matters because leading indicators support anticipation, while lagging indicators support confirmation. The planner specifies both: leading indicators to trigger plan revisions and contingency responses early, and lagging indicators to confirm that a revision or response was warranted. The leading indicators are the more valuable and the more contested, because their interpretation requires judgment about whether a movement is signal or noise.

The indicators enter the monitoring system treated in Chapter 20 of the Planning Manual. Each indicator is assigned a source, a frequency, a threshold, and a reporting line. The thresholds are derived from the scenario narratives: the values at which the indicator's movement is sufficient to conclude that the corresponding scenario is emerging. The reporting line specifies who is notified when a threshold is crossed, and what action the notification is to trigger, so that monitoring is an active input to plan revision rather than a passive record. The integration requires that the monitoring function and the foresight function agree on the indicator specifications, because the specifications carry the scenario logic into operational practice.

The indicator set is reviewed periodically, because the relationships among indicators, scenarios, and the system under study change over time. Some indicators lose their discriminating power as the system evolves; new indicators become available as data systems develop. The review is conducted jointly by the foresight function and the monitoring function, because neither has the full knowledge required: the foresight function knows the scenarios, and the monitoring function knows the data. Periodic review is what keeps the indicator set aligned with the scenarios it is meant to discriminate. The review cycle is timed to the scenario set's revision cycle, so that the indicators and the scenarios are revised in concert.

13.4 Scenarios and Plan Assumptions

Plan assumptions are the conditions the plan takes as given. Every plan rests on assumptions: about growth, about prices, about policy, about behaviour. The use of scenarios in assumption-setting is the disciplined selection of these conditions from the range of plausible futures, rather than their derivation from a single projection. The selection is documented, so the assumptions can be tested against the scenarios and revised when the scenarios change. An assumption that is not tied to a scenario is a private judgment; an assumption that is tied to a scenario is a position that can be defended, challenged, and updated.

The standard configuration, treated in Chapter 7 of the Planning Manual, is a central scenario surrounded by variant scenarios. The central scenario represents the planner's best assessment of the likely future; the plan's principal assumptions are drawn from the central scenario. The variant

scenarios represent futures that differ from the central case; the plan's contingency assumptions and trigger conditions are drawn from the variants. The distinction is functional: the central scenario supports the plan's main commitments, while the variants support the plan's flexibility. Where the scenario set is symmetric, with no central case, the planner selects a working central case for assumption-setting while preserving the symmetry for resilience testing.

Each assumption is documented with the scenario from which it is drawn, the reasoning that supports it, and the conditions under which it would need to be revised. An assumption drawn from the central scenario carries one set of revision conditions; an assumption drawn from a variant scenario carries another. The documentation is the basis on which the assumptions can be challenged, reviewed, and revised. A plan whose assumptions are not documented against the scenarios cannot be tested for resilience, because the conditions under which the assumptions hold have not been specified. The documentation also permits the assumptions to be carried forward across plan cycles, with revisions recorded against the reasons for them.

The practice of documenting assumptions against scenarios is established in the fiscal planning of several mature jurisdictions. The International Monetary Fund's Article IV consultations require member states to present fiscal assumptions against a central and a downside scenario. The United Kingdom Office for Budget Responsibility publishes its economic forecasts with a central case and scenario variants, with the assumptions tied to each; the United States Congressional Budget Office publishes its baseline alongside alternative scenarios, with the assumptions documented. The practice ensures that the assumptions are exposed to scrutiny and that the plan's vulnerability to alternative futures is visible. the doctrine treats this practice as the standard, not as a national choice.

13.5 Scenario-Based Contingency Planning

Scenario-based contingency planning is the use of the scenario set to design the responses that would be triggered if a variant scenario emerges. Contingency planning is treated in the Planning Manual at the general level; the assessment develops the specialty craft of deriving contingencies from scenarios. The scenario set provides the structure within which contingencies are designed: each variant scenario specifies the conditions that would require a response, and the response is designed to address those conditions. The output is a set of contingency plans, each tied to a scenario and to the indicators that signal its emergence. The contingency plans are stored alongside the main plan, ready for activation, and are reviewed whenever the scenario set is revised.

Each contingency plan has trigger conditions: the observable states at which the contingency is activated. The trigger conditions are derived from the scenario indicators treated in section 13.3: the thresholds at which the indicators' movement is sufficient to conclude that the corresponding scenario is emerging. The triggers are specified with the action they activate, the authority that activates it, and the timeline within which activation must occur. The specification is what distinguishes a contingency plan from a scenario description: a scenario describes what could happen, while a contingency plan specifies what will be done if it does. A trigger without an assigned authority is a signal without a response, and is not a contingency plan.

For each variant scenario, the planner designs a response that addresses the conditions the scenario describes. The response options may be classified by their lead time: pre-positioned responses that can be activated immediately, planned responses that require preparation, and contingent responses that require political authorisation. The classification governs the speed of response: pre-positioned responses are fastest, and contingent responses are slowest. The planner matches the response type to the scenario's lead time, which is the period between the indicator crossing the trigger and the scenario's conditions materialising. Where the lead time is short, the response must be pre-positioned; where the lead time is long, the response may be contingent, with the cost of preparation deferred until the trigger is crossed.

The Planning Manual treats contingency planning as a general discipline applicable across planning domains. the doctrine extends the doctrine by specifying how contingencies are derived from scenarios: the trigger conditions, the response options, and the scenario-to-contingency mapping. The extension is necessary because scenario-based contingency planning carries more structure than contingency planning in general: the scenarios provide a structured basis for the contingencies, and the indicator system provides a structured basis for the triggers. Where the Manual's general doctrine applies, the planner follows it; where the scenario specialty adds structure, the planner applies the doctrine's doctrine. The result is a contingency system that is integrated with the plan's assumptions, the plan's resilience profile, and the plan's monitoring system.

13.6 Chapter Summary

Scenario use is the application of constructed scenario sets to decisions. The use follows construction, and the construction record must be available to those who use the scenarios. The discipline is recognised across mature practice, from Shell's planning applications to the OECD's futures programme. The value of the scenario method is realised only when the set informs a decision, and the planner's responsibility is to make that translation explicit.

Testing plan resilience is the application of the scenario set to the plan to determine whether the plan succeeds under each plausible future. The test distinguishes robust plan elements, which succeed under every scenario, from scenario-dependent elements, which succeed under some and fail under others. Modification takes three forms: redesign for robustness, conversion to conditional commitments, and referral to contingency planning. The resilience profile is published alongside the plan, so that the political authority can see where the plan is robust and where it depends on conditions, and the profile is recorded in a form that supports re-testing when the plan or the scenarios are revised.

Scenario indicators are the observable measures that distinguish which scenario is emerging. Leading indicators support anticipation; lagging indicators support confirmation. The indicators are integrated into the monitoring system, with sources, frequencies, thresholds, and reporting lines. The indicator set is reviewed periodically to maintain its discriminating power as the system evolves, with the review timed to the scenario set's revision cycle.

Plan assumptions are the conditions the plan takes as given, selected from the range of plausible futures. The standard configuration is a central scenario surrounded by variants; the

central scenario supports the main commitments, and the variants support flexibility. Each assumption is documented with the scenario from which it is drawn and the conditions under which it would be revised. Scenario-based contingency planning uses the scenario set to design responses triggered by emerging scenarios, with trigger conditions derived from indicators and response options classified by lead time. The chapter extends the Planning Manual’s general contingency-planning doctrine by specifying the scenario-to-contingency mapping.

Chapter 14 — Expert Elicitation

Part IV: Foresight Methods — The structured methods that support the foresight discipline.

14.1 What Expert Elicitation Is

Expert elicitation is the structured collection of expert judgment on questions where data is insufficient for statistical inference. It draws on the knowledge of specialists — scientists, engineers, practitioners, policy officials — to characterise conditions that cannot be measured directly or projected from past observations. The method assembles judgment under protocols that make it transparent, reviewable, and reproducible. The output is a documented assessment, with the experts identified, the questions stated, the reasoning recorded, and the dissent preserved.

The distinction from informal consultation is structural. Informal consultation gathers advice through conversation or advisory committees without a fixed protocol, producing impressions that cannot be traced to their sources. Expert elicitation applies a defined method: the questions are written in advance, the experts are selected against criteria, the responses are captured in a standard form, and the synthesis follows auditable rules. Consultation may suffice for operational questions; elicitation is required when the assessment will enter a planning product and must withstand scrutiny. The line between the two is the protocol rather than the expertise.

The role in the foresight process is to supply judgment where evidence runs out. Horizon scanning produces signals but cannot characterise their significance without expert reading; driver identification requires experts to assess direction, strength, and time horizon; scenario construction depends on expert judgment to identify critical uncertainties and to test the internal consistency of narratives. The Planning Manual treats expert elicitation as an input to horizon scanning in Chapter 6; the doctrine develops the specialty craft across the principal methods and the management of their characteristic biases.

The RAND Corporation’s work from the 1950s onward is the principal comparative reference. RAND developed the Delphi method to forecast technological and strategic developments where statistical projection was unavailable, and the method’s protocols — anonymity, cyclic rounds, controlled feedback, structured aggregation — became the foundation of formal elicitation practice. Subsequent work by the United States Nuclear Regulatory Commission, the European Commission’s Joint Research Centre, and the United Kingdom Government Office for Science has refined the method, the selection of experts, and the synthesis. The doctrine that emerges is the doctrine treated here: elicitation is structured, documented, and bias-aware, and its results are presented as judgment under stated uncertainty rather than as fact.

14.2 The Delphi Method

The Delphi method is an iterative, anonymous, and structured elicitation conducted across multiple rounds of written questionnaires. Developed at the RAND Corporation in the 1950s for the forecasting of technological change, it has become the principal instrument for structured expert

elicitation where face-to-face interaction is impractical or undesirable. The defining features are anonymity among participants, the cyclic conduct of rounds, controlled feedback of previous responses, and structured aggregation of the final judgments. Anonymity suppresses the influence of status and personality, the rounds allow revision in light of the panel's reasoning, and controlled feedback prevents drift toward premature consensus.

The questionnaire is the instrument through which the method operates, and its design determines the quality of the elicitation. The questions are written in advance, in unambiguous language, with the response format specified — a numerical estimate, a probability assignment, a ranking, or a structured judgment. The questions are tested before the first round, to identify ambiguity and to confirm that the experts can answer them in the form requested. The instrument is the same for all experts and across all rounds, with revisions between rounds limited to clarification and to follow-up questions prompted by the previous round's responses.

The rounds are managed to converge the panel's judgment without forcing agreement. The first round elicits the experts' initial responses and reasoning; the second returns the aggregated responses — typically the median and the interquartile range — with a summary of the reasoning, and asks each expert to revise or to defend a position outside the interquartile range. Subsequent rounds repeat the cycle until the responses stabilise or until the marginal gain no longer justifies another round. Three rounds are standard, and more are rarely warranted. The facilitator does not push for consensus: persistent dissensus is recorded, because the dispersion of expert judgment is itself information about the depth of the uncertainty.

Convergence and dissensus are both documented as part of the elicitation's output. Convergence is reported as the central tendency of the final round — the median for numerical estimates, the modal response for categorical judgments — together with the dispersion around it. Dissensus is reported as the range of the final responses, the arguments advanced by the dissenting experts, and the issues on which the panel could not agree. The documentation includes the questionnaire, the experts' identities and affiliations, the response rates by round, the aggregated responses by round, and the reasoning captured.

14.3 Expert Panels

The expert panel is a face-to-face structured elicitation in which a convened group of experts assesses a defined question under a protocol that governs the discussion and the capture of input. The panel is the appropriate instrument where the question requires interaction among experts — where the resolution of disagreement depends on the direct exchange of reasoning, or where complexity demands that experts challenge each other's assumptions in real time. The panel differs from the Delphi method by convening the experts in the same room, with the gains of depth of exchange and the losses of exposure to group dynamics. The panel's protocols compensate for the losses through structured facilitation and documented dissent.

The selection of experts is the most consequential decision in the panel's design, governed by criteria stated in advance. The panel comprises experts with the relevant substantive knowledge, drawn from the disciplines and institutions that hold that knowledge, with diversity across disciplines, affiliations, and perspectives to guard against the shared blind spots of any single

community. The number is small — typically six to twelve — large enough to span the relevant perspectives and small enough to permit genuine discussion. The selection is documented, with the criteria, the candidates considered, and the reasons for the final choice, so the panel's composition can be reviewed.

The facilitation is conducted by a practitioner whose role is to manage the discussion according to the protocol, free of substantive contribution. The facilitator opens the panel, invites each expert in turn to present an initial position, manages the subsequent exchange to ensure that each expert is heard and that none dominates, and intervenes to surface disagreement rather than allow it to settle into false consensus. The facilitator distinguishes between substantive contributions, which are the experts' to make, and procedural interventions, which are the facilitator's. The facilitation's quality determines whether the panel produces a structured assessment or a conversation that converges on the most assertive voice.

The synthesis of panel input captures the points of agreement, the points of disagreement, and the reasoning on each. The synthesis is a structured record of where the experts concur, where they differ, and what each position rests on, without forcing a single position from the panel. It is prepared by the foresight function from the panel's record and is reviewed by the experts before it is finalised, to confirm that the synthesis represents the panel's views rather than the practitioner's. The reviewed synthesis is the panel's output, and it enters the foresight product with the dissent preserved alongside the agreement.

14.4 Structured Workshops

The structured workshop is a facilitated group elicitation that uses scenarios, driver mapping, or cross-impact analysis as the framework through which expert input is collected. The workshop differs from the expert panel by its method: the panel assesses a question through discussion, while the workshop produces an analytical product through a structured exercise in which the participants work together on a defined task. The workshop is the appropriate instrument where the foresight question requires the integration of expert judgment with structured analysis — where the participants are asked to construct, to map, or to test, as well as to assess. The Shell scenario workshops, conducted continuously since the early 1970s, are the principal comparative reference.

The workshop's design begins with the question the workshop is to address and the analytical product it is to produce. A scenario workshop produces a set of scenarios or tests the consistency of scenarios already drafted; a driver mapping workshop produces a driver map with the relationships among the drivers; a cross-impact workshop produces a matrix that captures the experts' judgment of how each driver affects each other. The design specifies the participants, the materials, the sequence of exercises, and the time allocated to each. The design is documented in advance, reviewed by the commissioning authority, and distributed as the workshop's terms of reference.

The facilitation is conducted by one or more practitioners whose role is to guide the participants through the exercises, to capture the input as it is produced, and to manage the group dynamics that arise when experts work together under time pressure. The facilitator presents each exercise, manages the discussion to keep it on the exercise, captures the output on the worksheets or the shared display, and intervenes when the discussion stalls or when a participant dominates.

The facilitator refrains from substantive contribution and from editorial comment, so the workshop's output remains the participants' rather than the practitioner's.

The documentation of outputs is conducted in parallel with the workshop, rather than deferred to a later stage. Each exercise's output is captured as it is produced — on worksheets, on a shared display, in the facilitator's notes — and is reviewed with the participants at the end of each exercise to confirm that the capture is accurate. The workshop's record comprises the captured outputs, the participants' written submissions where these are called for, the facilitator's notes, and the materials distributed in advance. The product is prepared from the record by the foresight function, is reviewed by the participants, and is then issued under the commissioning authority's responsibility.

14.5 Managing Expert Bias

Expert judgment is subject to systematic biases that, unmanaged, distort the elicitation's results. The biases are documented in the cognitive psychology literature, drawing principally on the work of Daniel Kahneman, Amos Tversky, and their successors, and they operate on experts as they operate on other informed observers. The principal biases in elicitation practice are availability bias, anchoring, overconfidence, and groupthink, and each affects the elicitation through a distinct mechanism. The biases cannot be eliminated, because they are features of judgment under uncertainty; the doctrine is to manage them through method design, expert selection, and structured synthesis, so their effect is reduced and the residual effect is documented.

Availability bias is the tendency to judge an event's likelihood by the ease with which examples come to mind, which leads experts to over-weight recent, vivid, or personally experienced cases. The bias is managed by structuring the question to require consideration of the full range of evidence, by asking the expert to identify the base rate before estimating the specific case, and by aggregating judgments across experts whose experiences differ. Anchoring is the tendency to adjust insufficiently from an initial value, which leads experts' estimates to be pulled toward any number placed before them. The bias is managed by withholding aggregate values until each expert has recorded an independent judgment and by separating the elicitation of the estimate from the elicitation of the reasoning.

Overconfidence is the tendency to express greater certainty than the evidence warrants, which leads experts to produce estimates with intervals too narrow to contain the true value and to assign probabilities that are too extreme. The bias is managed by requiring experts to justify their confidence levels, by asking for ranges rather than point estimates, and by feeding back the calibration of previous elicitations where track records are available. Groupthink is the tendency of convened groups to converge on a shared position prematurely, suppressing dissent in the interest of harmony or under the influence of a dominant voice. The bias is managed by the Delphi method's anonymity, by the panel facilitator's solicitation of dissent, and by the workshop's structured exercises that require individual input before group discussion.

The management of bias through method design, expert selection, and structured synthesis is the doctrine's three-part response. Method design builds the bias controls into the elicitation protocol: anonymity, the cyclic conduct of rounds, the structure of the question, the order of the

prompts, and the capture of reasoning alongside the estimate. Expert selection reduces the shared biases of any single community by drawing experts from different disciplines, institutions, and perspectives, and by documenting the diversity so the residual shared bias can be assessed. Structured synthesis addresses the residual biases by aggregating across experts, preserving dissent, and reporting the dispersion alongside the central tendency.

14.6 Synthesising Expert Input

The synthesis of expert input is the stage at which the elicitation's raw material — the experts' individual judgments and reasoning — is converted into a foresight product that can enter the planning process. The synthesis is the foresight function's responsibility, conducted under the commissioning authority's direction, and governed by the same standards of documentation and transparency as the elicitation itself. The synthesis is a structured assessment of what the elicitation establishes, what it leaves uncertain, and what the planning process should do with both, distinct from a summary of the experts' statements. The discipline of the synthesis determines whether the elicitation's investment is realised in the planning product.

The aggregation of judgments follows rules stated in advance and applied consistently. For numerical estimates, the standard aggregation is the median of the experts' final-round responses, with the interquartile range reported as the measure of dispersion; the median is preferred to the mean because it is robust to outliers and because it represents the central tendency of a small panel without imposing a distributional assumption. For categorical judgments, the aggregation reports the distribution of responses and the modal category, with the dissenting positions preserved. For qualitative judgments, the aggregation identifies the points of agreement, the points of disagreement, and the reasoning on each, without forcing a single position.

The treatment of dissensus is the synthesis's most demanding element. Dissensus is information about the depth of the uncertainty and the presence of legitimate disagreement among informed experts, treated as a finding rather than a failure. The synthesis reports the dissensus explicitly: the range of positions, the experts who hold each, and the arguments advanced. Where the dissensus reflects different assumptions about the underlying processes, the synthesis identifies the assumptions and their consequences; where it reflects genuine uncertainty, the synthesis reports the uncertainty and its implications for the plan. The synthesis preserves the dissensus rather than collapsing it into a single position, because the dissensus is itself the input the plan must address.

The documentation of the basis for the synthesis is the condition under which the synthesis can be reviewed, challenged, and revised. The documentation includes the elicitation protocol, the experts' identities and affiliations, the questions, the responses by round, the aggregation rules, the reasoning captured, and the practitioner's assessment of what the elicitation establishes and leaves uncertain. The documentation is sufficient for a qualified reviewer to follow the synthesis from the experts' judgments to the foresight product, to identify where the practitioner's judgment entered, and to test whether the synthesis follows from the elicitation. The documentation is the synthesis's defence against the charge that it is the practitioner's opinion dressed as the experts' judgment.

14.7 Chapter Summary

Expert elicitation is the structured collection of expert judgment on questions where data is insufficient for statistical inference, conducted under protocols that make the judgment transparent, reviewable, and reproducible. The method is distinguished from informal consultation by its protocol, and the protocol is the line that determines whether the assessment can enter a planning product. The role in the foresight process is to supply judgment where evidence runs out — in horizon scanning, driver identification, scenario construction, and risk assessment. The RAND Corporation's tradition from the 1950s is the principal comparative reference, refined by subsequent practice across mature jurisdictions.

The principal methods are the Delphi method, the expert panel, and the structured workshop. The Delphi method is iterative, anonymous, and conducted across rounds of written questionnaires, with convergence and dissensus both documented. The expert panel is a face-to-face structured elicitation, governed by facilitation and the synthesis of panel input. The structured workshop is a facilitated group elicitation using scenarios, driver mapping, or cross-impact analysis as the framework, with the Shell scenario workshops as the principal comparative reference.

Expert judgment is subject to systematic biases — availability bias, anchoring, overconfidence, and groupthink — that the elicitation's protocols are designed to manage. The doctrine is three-part: method design builds the controls into the protocol; expert selection reduces shared biases through diversity; structured synthesis addresses the residual biases through aggregation, the preservation of dissent, and the disclosure of dispersion. The doctrine promises judgment whose biases have been identified, managed, and disclosed, rather than judgment free of bias.

The synthesis of expert input converts the elicitation's raw material into a foresight product, under rules of aggregation stated in advance and applied consistently, with dissensus preserved rather than collapsed, and with the basis for the synthesis documented to the standard of reviewability. The discipline of the synthesis distinguishes a foresight product from an opinion piece, and it is the condition under which the elicitation's investment is realised in the plan. the doctrine extends the Planning Manual's general account with the specialty craft of method selection, bias management, and synthesis, and the chapters that follow develop the foresight function's other methods in the same depth.

Chapter 15 — Cross-Impact Analysis

The assessment of relationships among drivers and trends.

15.1 What Cross-Impact Analysis Is

Cross-impact analysis is the structured assessment of the relationships among drivers and trends. The method asks, for each pair of drivers, what effect a change in one would have on the other. The answer is recorded as a pairwise judgment, and the assembly of all such judgments produces a structured map of the driver network. The method's purpose is to make the relationships among drivers explicit, documented, and subject to review, rather than left implicit in the analyst's working understanding. Cross-impact analysis does not produce a forecast; it produces the relational map the planner uses when constructing scenarios and assessing strategic risk.

Cross-impact analysis extends the driver mapping treated in Chapter 10. Driver mapping identifies the drivers and assesses each driver's direction, strength, and time horizon; cross-impact analysis assesses the relationships among them. The two are conducted in sequence, with the driver map as the input and the cross-impact matrix as the analytical product that converts the map from a list into a network. The Planning Manual's Chapter 8 sets out the method at the general level; the doctrine develops the specialty craft of matrix construction, leverage-point identification, feedback-loop analysis, and scenario consistency testing.

The method's principal outputs are two. The first is the identification of leverage points: drivers that have high influence on others, such that a change in them produces a cascade of effects across the network. The second is the identification of feedback loops: circuits of relationships through which a change in one driver feeds back, after passing through other drivers, to amplify or dampen the original change. Both outputs are essential to scenario construction: leverage points determine which drivers the scenarios should foreground, and feedback loops determine whether a scenario's internal dynamics are consistent.

The method rests on a discipline of explicit judgment. Each cell in the cross-impact matrix is the analyst's assessment, supported by evidence and reasoning, of the effect of one driver on another. The assessment is recorded, so that it can be reviewed, challenged, and revised; the recording is what permits the matrix to be maintained across staff turnover and to be audited by readers who were not present at its construction. A matrix assembled in a workshop and not recorded has no institutional life beyond the room in which it was produced; the documentation is what converts the workshop's proceedings into an analytical product that survives the participants' departure.

15.2 The Cross-Impact Matrix

The cross-impact matrix is the structured representation of the pairwise assessments. The matrix has the drivers as both rows and columns, so that each cell records the assessed impact of the row driver on the column driver; the diagonal cells are left blank, since a driver's impact on

itself is not a cross-impact. The matrix is square, with as many rows and columns as there are drivers, and a matrix with fifteen drivers contains 210 assessed cells while one with thirty contains 870. The size is a design choice, governed by the trade-off between analytical completeness and analytical tractability, and exercises of more than thirty drivers are usually partitioned into sub-matrices to preserve tractability.

The scale of assessment is defined before the matrix is constructed, and is applied consistently across all cells. A common scale is five points: strongly negative, weakly negative, none, weakly positive, and strongly positive, where the polarity refers to the direction of the effect on the column driver. Some exercises use a seven-point scale; some use a three-point scale; the choice is governed by the analysts' ability to discriminate consistently. The scale's anchors are defined in writing, with criteria for each level, so that two analysts assessing the same cell would arrive at the same rating. The discipline of the anchors is what permits the matrix to be assembled by a team and to be combined across exercises.

The matrix is constructed through a combination of expert elicitation and structured workshop. The drivers are drawn from the driver map, refined to the set the cross-impact exercise will consider. Each cell is then assessed: the analyst or the workshop participants judge the impact of the row driver on the column driver, applying the defined scale and the defined criteria. The assessment is conducted cell by cell, with the reasoning recorded for each non-zero rating; the reasoning is what distinguishes a defensible matrix from a table of opinions. Where the analysts disagree, the disagreement is recorded, and the cell carries the assessed rating together with the documented dissent.

The matrix is documented to a standard that allows it to be reviewed and revised. Each cell's rating is recorded with its basis: the evidence, the reasoning, and the source of the judgment. The matrix is dated, attributed to the analysts who produced it, and stored in a form that allows comparison with subsequent revisions. The documentation is the condition under which the matrix can be audited: a reader who was not present at its construction must be able to reconstruct the basis for each rating, to challenge the ratings, and to revise them as new information arrives. A matrix that is not documented to this standard is a private working paper, and it loses its value at the next staff transition.

15.3 Identifying Leverage Points

Leverage points are drivers that have high influence on others. They are identified from the cross-impact matrix by summing, for each row, the absolute values of the cells in that row: the row total measures the row driver's influence on the network as a whole. A driver whose row total is high exerts strong influence on many other drivers, and a change in that driver produces a cascade of effects across the network. A driver whose row total is low exerts little influence, and a change in that driver produces few second-order effects. The row totals are ranked, and the drivers at the top of the ranking are the leverage points.

The complementary measure is the column total, which sums the absolute values of the cells in each column. The column total measures the degree to which the column driver is influenced by the network: a driver with a high column total is heavily influenced by other drivers, and its

trajectory is largely determined by forces outside its own development. A driver with a low column total is weakly influenced by others, and its trajectory is determined by its own dynamics. The drivers with high column totals and low row totals are indicators: their changes signal that changes in the influencing drivers have occurred, even before those changes are directly observed.

The distinction between leverage points and indicators is the distinction between drivers the plan should act on and drivers the plan should monitor. A leverage point is a driver where intervention can produce wide effects; an indicator is a driver where observation can detect change produced elsewhere in the network. The plan's strategy addresses leverage points, because they are the points at which the plan's leverage is greatest; the plan's monitoring tracks indicators, because they are the points at which the plan's information is most efficiently gathered. The cross-impact analysis provides the analytical basis for both decisions, and the discipline of recording the matrix ensures that the basis is available for review.

The use of leverage points in scenario construction follows. The scenarios differ along the critical uncertainties, as treated in Chapter 9; the leverage points are the drivers whose development the scenarios must represent in detail, because their development shapes the development of many others. A scenario that resolves a leverage point one way produces a markedly different future from a scenario that resolves the same leverage point the other way. The leverage points are also the drivers the scenarios stress-test: the scenario set is constructed so that the leverage points are exercised across their plausible range, and the implications of each resolution are traced through the network the matrix represents.

15.4 Identifying Feedback Loops

A feedback loop is a circuit of relationships through which a change in one driver feeds back, after passing through other drivers, to the original driver. Reinforcing loops amplify the original change: a change in driver A increases driver B, which increases driver C, which increases driver A further. Balancing loops dampen the original change: a change in driver A increases driver B, which decreases driver C, which decreases driver A back toward its original level. The two types produce qualitatively different system behaviour, and the identification of the loops in the network is essential to understanding how the system will behave under stress.

Feedback loops are identified from the cross-impact matrix by tracing paths through the cells. A path is a sequence of drivers, each linked to the next by a non-zero cell, that returns to the starting driver. The path is a reinforcing loop if the product of the signs along the path is positive, and a balancing loop if the product is negative. The tracing is straightforward for short loops but becomes difficult for longer loops, because the number of paths grows rapidly with the matrix's size. Software tools support the tracing for matrices of more than a dozen drivers, and the mature practice uses them where the manual tracing would be impractical.

The significance of feedback loops for understanding system behaviour is large. A system dominated by reinforcing loops exhibits accelerating change: a small perturbation grows over time, and the system moves away from its prior state. A system dominated by balancing loops exhibits stabilising change: a perturbation is dampened, and the system returns toward its prior state. The dominance is determined by the loops' strengths as well as their presence, and the matrix's ratings

provide the basis for assessing the relative strengths. A scenario constructed for a system dominated by reinforcing loops must allow for acceleration; a scenario constructed for a system dominated by balancing loops must allow for stability.

The use of feedback loops in scenario consistency testing follows. A scenario that contains a reinforcing loop must account for the loop's effects: the loop will amplify the changes the scenario posits, and the scenario's outcomes will tend toward the extreme unless the loop is bounded by a balancing mechanism. A scenario that contains a balancing loop must account for the loop's effects: the loop will dampen the changes the scenario posits, and the scenario's outcomes will tend toward moderation unless the loop is overridden by a stronger reinforcing mechanism. The feedback-loop analysis is the principal test of a scenario's internal dynamics, and a scenario that posits changes incompatible with its loops is internally inconsistent and must be revised.

15.5 Using Cross-Impact in Scenario Construction

Cross-impact analysis tests the internal consistency of scenarios by checking whether the combinations of driver outcomes the scenario posits are compatible with the relationships the matrix records. A scenario is a combination of outcomes for the critical uncertainties and the leverage points, set against the predetermined elements; the cross-impact matrix records the relationships among the drivers, and the consistency test asks whether the posited combination is supported or contradicted by those relationships. The test is conducted by tracing, for each pair of posited outcomes, the implication of the matrix's cell. A pair that the matrix rates as strongly reinforcing is consistent if the scenario posits both outcomes in the same direction; a pair that the matrix rates as strongly opposing is consistent only if the scenario posits the outcomes in opposite directions.

The identification of incompatible driver combinations is the principal output of the consistency test. An incompatible combination is a pair, or a set, of posited outcomes that the matrix's relationships cannot support: a scenario that posits a strong increase in driver A and a strong decrease in driver B, where the matrix records a strong reinforcing relationship between A and B, is incompatible, because the reinforcing relationship implies that A and B move together. The incompatibility may be partial — the combination is unlikely but not impossible — or complete — the combination is inconsistent with the matrix and cannot stand. The test identifies the incompatibilities, classifies them, and records them as items the scenario must address.

The revision of scenarios to achieve consistency follows. A scenario with an incompatibility is revised by altering one or more of the posited outcomes, by introducing a mediating driver that resolves the incompatibility, or by relaxing the scenario's assumptions. The revision is governed by the matrix: the alteration must produce a combination the matrix supports, and the mediating driver must be one the matrix recognises. A scenario that cannot be revised to consistency is abandoned, and the abandonment is recorded as a finding of the cross-impact analysis: the combination the scenario posited is not plausible, and the planning function should not prepare for it. The recording of abandoned scenarios is part of the analysis's contribution, because it narrows the range of futures the plan must address.

Cross-impact analysis also informs the selection of the scenario set. The scenario set spans the range of plausible futures, and the range is defined by the critical uncertainties' plausible combinations. The cross-impact matrix constrains the plausible combinations, by identifying the combinations the matrix supports and ruling out the combinations the matrix contradicts. A scenario set constructed without the matrix's constraint may include combinations that are internally inconsistent; a scenario set constructed with the matrix's constraint is grounded in the network's structure. The discipline of the consistency test is the condition under which the scenario set is a defensible representation of the plausible futures, rather than a collection of narratives the analysts found interesting.

15.6 Chapter Summary

Cross-impact analysis is the structured assessment of the relationships among drivers and trends. The method converts the driver map from a list into a network, by assessing each pairwise relationship and recording it in a matrix. The method's principal outputs are the identification of leverage points — the drivers with high influence on others — and the identification of feedback loops — the circuits through which change is amplified or dampened. The method extends the doctrine set out in Chapter 8 of the Planning Manual; the doctrine develops the specialty craft of matrix construction, leverage-point identification, feedback-loop analysis, and scenario consistency testing.

The cross-impact matrix is the method's principal instrument. The matrix has the drivers as both rows and columns, with each cell recording the assessed impact of the row driver on the column driver on a defined scale. The matrix is constructed through expert elicitation and structured workshop, with the reasoning for each non-zero rating recorded. The documentation allows the matrix to be reviewed, challenged, and revised, and is the condition under which the matrix is an institutional product rather than a private working paper.

Leverage points are identified from the matrix's row totals, and indicators from the column totals. Leverage points are the drivers the plan should address, because intervention at those points produces the widest effects; indicators are the drivers the plan should monitor, because their changes signal developments elsewhere in the network. Feedback loops are identified by tracing paths through the matrix, and are classified as reinforcing or balancing by the product of the signs along the path. The loops' dominance determines whether the system exhibits accelerating or stabilising behaviour, and the scenarios must be constructed to be compatible with the loops they contain.

Cross-impact analysis tests the internal consistency of scenarios by checking whether the posited combinations of driver outcomes are supported by the matrix's relationships. Incompatible combinations are identified, classified, and addressed through revision or abandonment. The matrix also constrains the selection of the scenario set, by ruling out combinations the matrix contradicts. The discipline of the consistency test is the condition under which the scenario set is a defensible representation of the plausible futures, and the chapters that follow treat the use of cross-impact analysis in strategic risk assessment and in the design of monitoring systems.

Chapter 16 — Trend Analysis and Extrapolation

The identification, measurement, and projection of trends over the planning horizon.

16.1 What Trend Analysis Is

Trend analysis is the systematic identification, measurement, and projection of trends over the planning horizon. A trend is a direction of change observed over time in a measurable quantity; trend analysis converts that observation into a structured assessment the planning function can use. The method produces three outputs in sequence: an identified set of trends relevant to the planning question; a measurement of each trend in a form that supports projection; and a projection of each trend over the horizon, with an assessed uncertainty. The method belongs to the foresight discipline because it converts observed change into a structured input to the plan's assumptions, and its outputs feed the baseline treated in Chapter 9 of the Planning Manual.

Trend analysis is paired with driver mapping in the Planning Manual's doctrine, and the two are distinct operations. A trend is what is observed; a driver is what causes the observation. Population ageing is a trend; declining fertility and increasing life expectancy are the drivers. Trend analysis works on the observation, measuring its trajectory and projecting it forward; driver mapping, treated in Chapter 10, works on the causes, assessing their direction, strength, and interaction. A trend projected without an understanding of its drivers is an extrapolation whose continuation cannot be assessed, and a driver mapped without measurement of the trend it produces is an analysis whose empirical consequence is unknown.

Trend analysis is conducted in the diagnostic phase of the foresight process, after the planning question has been defined and before the scenario construction begins. The trends identified, measured, and projected in this phase become the empirical base on which the scenarios are constructed and against which the plan's assumptions are tested. The method's position in the process gives it a particular character: its outputs must be defensible on their own terms, because they will be cited in the plan and scrutinised by the strategy function, the political authority, and the audit bodies that review the plan. The discipline of documentation is therefore as important as the discipline of analysis, and the doctrine treats both.

16.2 Trend Identification

Trend identification begins with the planning question and works outward to the trends that bear on it. A planning question is a statement of what the plan must decide, framed in operational terms; the trends relevant to it are those whose trajectory would materially affect the decision. For a national development plan, the relevant trends include demographic, economic, fiscal, social, environmental, and technological trends; for a sectoral plan, the relevant trends are the sectoral trends plus the broader trends that shape the sector's environment. The identification is selective: the planning function cannot analyse every measurable trend, and the discipline of selection is to confine the analysis to the trends whose trajectory would change the plan's commitments.

The sources of trend identification are statistical data, administrative data, survey data, and expert observation, used in combination. Statistical data are produced by the statistics office under the standards set out in the International Monetary Fund's Data Quality Assessment Framework and the United Nations' Fundamental Principles of Official Statistics, and these are the authoritative source for demographic, economic, and social trends. Administrative data are produced by government agencies in the course of their operations — tax records, health records, education enrolments, social benefit rolls — and cover trends the statistical system does not capture directly. Survey data, where conducted to methodological standard, fill gaps the administrative record does not cover, such as household income, time use, and subjective wellbeing. Expert observation, drawn from sectoral specialists and from the international literature, identifies trends that are not yet visible in the data but are recognised by practitioners as significant.

The sources are cross-validated wherever the planning question permits. A demographic trend identified from the statistics office's population projections is checked against the administrative records of births, deaths, and migration; an economic trend identified from the national accounts is checked against the tax records and the labour force survey. The cross-validation serves two purposes: it confirms that the trend is real and not an artefact of a single source, and it reveals the magnitude of the uncertainty in the trend's measurement. Where the sources diverge, the divergence is investigated and recorded; the resolution is documented so that subsequent users of the trend analysis can reconstruct the basis on which the divergence was resolved.

Each identified trend is documented in a trend register, which records the trend's name, definition, source, temporal coverage, and relevance to the planning question. The register is the planning function's working inventory of the trends it will measure and project, and it is the basis on which the analysis is reviewed and revised. The register is maintained across planning cycles, so that trends identified in one cycle are available for re-assessment in the next; identifying trends anew for each cycle wastes the prior analysis and loses the institutional memory of how the trends were understood. The register's discipline is the condition under which the trend analysis is an institutional product rather than a personal working paper of the analysts who happened to be present.

16.3 Trend Measurement

Trend measurement applies the data identified in the identification phase to produce a quantitative characterisation of each trend. The characterisation has three elements: the level of the trend at the present, the rate of change over the historical period, and the pattern of change — whether the rate is constant, accelerating, decelerating, or oscillating. These three elements together describe the trend's trajectory, and they are the basis on which the projection will be constructed. The measurement is conducted on a time series of sufficient length to distinguish the trend from the cyclical and seasonal variation that overlays it; the IMF's Quarterly National Accounts Manual and the OECD's compilation guides set out the decomposition methods the planning function applies.

The data quality requirements are those the IMF's Data Quality Assessment Framework sets out: accuracy, reliability, periodicity, timeliness, and consistency. Accuracy is the degree to which the data measure what they purport to measure; reliability is the degree to which successive

measurements of the same phenomenon produce the same result; periodicity is the frequency with which the data are produced; timeliness is the lag between the reference period and the publication; consistency is the degree to which the data are comparable across periods and across datasets. The planning function assesses each source against these dimensions before using it, and records the assessment in the trend register. Where the data fall short on one or more dimensions, the measurement is qualified accordingly, and the qualification is carried into the projection and the uncertainty assessment.

The measurement is conducted at the level of disaggregation the planning question requires. A demographic trend is measured at the national level and at the sub-national level, by age group, by sex, and by migration status, because the plan's commitments may apply differently to each sub-population. An economic trend is measured at the aggregate level and by sector, by region, and by firm size, because the plan's interventions are typically targeted at sub-aggregates rather than the whole. The disaggregation is governed by the planning question and by the data's capacity to support finer breakdowns; disaggregation beyond the data's capacity produces spurious precision and is avoided. The principle is that the measurement should be at the level the plan will act at, and that the disaggregation should be no finer than the data can support.

The time series on which the measurement is based is documented to a standard that allows the measurement to be reproduced. The series' source, the adjustments applied, the revisions history, and the level at which the series is published are all recorded; the adjustments include seasonal adjustment, rebasing, and chain-linking, each applied on the basis of the methodology the statistics office publishes. The documentation is the condition under which the measurement can be audited, revised, and updated in subsequent cycles; a measurement whose basis cannot be reconstructed loses its value when the analyst who made it leaves the function. The documentation also records the measurement's limitations, so that the projection phase can carry those limitations forward into the uncertainty assessment rather than treating the measurement as if it were exact.

16.4 Trend Projection Methods

Trend projection extends the measured trajectory over the planning horizon. Three methods are used, alone or in combination: simple extrapolation, model-based projection, and expert-based projection. Simple extrapolation extends the observed rate of change forward, on the assumption that the conditions producing the trend will continue; the method is suited to trends whose drivers are stable and whose measurement is precise, such as demographic trends over a five-year horizon. Model-based projection uses a structural model that relates the trend to its drivers and projects the trend by projecting the drivers; the method is suited to trends whose drivers are themselves projected, such as economic trends projected through a macroeconomic model. Expert-based projection uses structured expert elicitation to produce a projection where neither extrapolation nor modelling is defensible; the method is suited to trends that are newly emergent, poorly measured, or subject to discontinuities the historical record does not capture.

The choice of method is governed by the trend's characteristics, the data's quality, and the horizon's length. Extrapolation is preferred where the trend is well-established, the data are of high quality, and the horizon is short relative to the trend's dynamics; the method's strength is its

transparency, and its weakness is its inability to detect turning points before they occur. Model-based projection is preferred where the trend's drivers are well-understood and where models of acceptable quality are available; the method's strength is its capacity to capture the structural relationships that govern the trend, and its weakness is its dependence on the model's assumptions, which are themselves uncertain. Expert-based projection is preferred where the other methods cannot be applied defensibly; the method's strength is its capacity to incorporate information the data do not yet contain, and its weakness is its dependence on the experts' judgment, which must be elicited, documented, and reviewed with the same rigour as a statistical method.

Every projection is accompanied by an assessment of its uncertainty. The assessment takes one of two forms: a quantitative confidence interval, where the method produces one, or a qualitative assessment of the projection's reliability, where it does not. Quantitative confidence intervals are produced by the statistical methods — extrapolation produces intervals from the historical variation, and model-based projection produces intervals from the model's error structure — and they are stated with the projection in the documentation. Qualitative assessments classify the projection's reliability as high, medium, or low, on defined criteria: the data quality, the trend's stability, the method's appropriateness, and the horizon's length. The uncertainty assessment is the condition under which the projection's users can judge how much weight to place on the projection and how much contingency to allow for its revision.

The projection is documented to a standard that allows it to be reviewed and revised. The documentation records the method, the inputs, the assumptions, the uncertainty assessment, and the analysts responsible; it records also the alternative projections considered and rejected, with the reasons for the rejection. The documentation is the condition under which the projection is auditable: a reader who was not present at its construction must be able to reconstruct how it was produced, to challenge its assumptions, and to revise it as new information arrives. The discipline of documentation is what distinguishes a projection that has institutional standing from a projection that is the personal view of the analyst; the former survives staff transitions and planning cycles, and the latter does not.

16.5 Assessing Trend Stability

Trend stability is the assessed likelihood that the trend will continue on its projected trajectory over the planning horizon. A stable trend is one whose drivers are well-understood, whose measurement is precise, and whose dynamics are unlikely to be disrupted within the horizon; an unstable trend is one whose drivers are changing, whose measurement is imprecise, or whose dynamics are subject to discontinuities the historical record does not capture. The assessment is qualitative, drawing on the quantitative evidence produced in the measurement and projection phases, and it is the bridge between the trend analysis and the scenario construction that follows. The assessment is recorded in the trend register, alongside the measurement and the projection, so that the basis for the stability judgment is open to review.

The indicators of stability are four. The first is the trend's historical persistence: a trend that has held its direction and rate over a long period is more likely to continue than one that has been established only recently. The second is the stability of the drivers: a trend whose drivers are

themselves stable — demographic drivers, for instance, change slowly — is more likely to continue than one whose drivers are volatile. The third is the presence of stabilising mechanisms: a trend governed by balancing feedback loops, identified in the cross-impact analysis treated in Chapter 15, is more likely to continue than one governed by reinforcing loops that may accelerate to a turning point. The fourth is the absence of anticipated disruptions: a trend for which no major policy change, technological shift, or environmental shock is anticipated within the horizon is more likely to continue than one for which such disruptions are anticipated.

Unstable trends are treated differently from stable trends in the foresight process. A stable trend is treated as a predetermined element: it is held constant across the scenarios, and the scenarios explore the uncertainties that lie on top of it. An unstable trend is treated as a critical uncertainty: its trajectory varies across the scenarios, and the scenario set is constructed to span its plausible range. The distinction is consequential because the scenario set's structure depends on it: too many unstable trends included as predetermined elements produces scenarios that are unrealistically close together, and too many stable trends included as critical uncertainties produces scenarios that are unrealistically far apart. The stability assessment is therefore the principal input to the scenario construction's design, and it is conducted with the rigour that responsibility requires.

Stability is re-assessed at each planning cycle, because trends that were stable in one cycle may become unstable in the next. A demographic trend that was stable for decades may become unstable if migration policy changes; an energy demand trend that was stable may become unstable if a technological shift alters the energy intensity of output. The re-assessment draws on the trend register, which permits the comparison of the current cycle's assessment with the prior cycle's, and on the cross-impact analysis, which permits the assessment of changes in the driver network. The re-assessment is the condition under which the trend analysis remains current: a trend analysis that is not re-assessed becomes stale, and the plan that rests on a stale analysis is exposed to the changes the analysis failed to detect.

16.6 The Relationship to the Baseline

Trend analysis is the principal empirical input to the baseline treated in Chapter 9 of the Planning Manual. The baseline is the central projection of the planning environment against which the plan's commitments are tested; it comprises the projected trajectories of the trends the plan treats as predetermined, combined into an integrated picture of the future the plan assumes. The trend analysis produces the individual projections; the baseline construction integrates them, reconciling their interactions and ensuring their consistency. The relationship is sequential: trend analysis first, baseline construction second, with the baseline drawing on the trend register and on the projections the trend analysis has documented.

The trend projection and the baseline are distinct products. A trend projection is the projected trajectory of a single trend, with its uncertainty assessment; it is the output of the trend analysis. The baseline is the integrated projection of the planning environment, constructed from the trend projections and from the assumptions the plan makes about the trends' interactions; it is the output of the baseline construction. The distinction matters because the two products serve different

purposes: the trend projection is an input to the analysis, used by the sectoral planners and the scenario constructors, while the baseline is the planning reference, used by the strategy function and the political authority. Conflating the two — treating a single trend projection as if it were the baseline, or constructing the baseline without explicit trend projections — produces analytical products that are neither defensible as projections nor usable as baselines.

The integration of trend analysis into baseline construction follows a defined sequence. The trend projections are assembled from the trend register, with their uncertainty assessments carried forward, and checked for consistency using the cross-impact matrix to identify pairs of projections whose combination the matrix contradicts; inconsistent combinations are resolved by revising one or both projections, with the revision documented. The consistent set of projections is then combined into the baseline, with the uncertainty assessments aggregated to produce a baseline uncertainty range. The baseline is documented to the same standard as the trend projections, and it is reviewed and revised at each planning cycle, drawing on the re-assessed trend analysis. The discipline of the integration is what gives the baseline its analytical standing, and a baseline constructed without that discipline is a projection presented as a reference rather than a reference constructed on defensible analytical foundations.

16.7 Chapter Summary

Trend analysis is the systematic identification, measurement, and projection of trends over the planning horizon. The method produces three outputs in sequence: an identified set of trends relevant to the planning question; a measurement of each trend in a form that supports projection; and a projection of each trend over the horizon, with an assessed uncertainty. The method is paired with driver mapping, treated in Chapter 10, and the two are complementary: trend analysis works on the observation, driver mapping on the causes. The method extends the doctrine set out in Chapter 8 of the Planning Manual, going deeper on the specialty craft of identification, measurement, projection, stability assessment, and integration into the baseline.

Trend identification draws on statistical data, administrative data, survey data, and expert observation, with sources cross-validated wherever the planning question permits. Each identified trend is documented in a trend register, which is maintained across planning cycles. Trend measurement applies the data to produce a quantitative characterisation of each trend’s level, rate, and pattern, conducted at the level of disaggregation the planning question requires, and on a time series of sufficient length to distinguish the trend from cyclical and seasonal variation. The data quality requirements are those the IMF’s Data Quality Assessment Framework sets out, and the measurement’s limitations are carried forward into the uncertainty assessment.

Trend projection uses simple extrapolation, model-based projection, or expert-based projection, with the method chosen for the trend’s characteristics, the data’s quality, and the horizon’s length. Every projection is accompanied by an uncertainty assessment, quantitative where the method produces one and qualitative where it does not. Trend stability is assessed on four indicators — historical persistence, driver stability, presence of stabilising mechanisms, and absence of anticipated disruptions — and the assessment determines whether the trend is treated as a predetermined element or as a critical uncertainty in the scenario construction. Trend analysis is

the principal empirical input to the baseline treated in Chapter 9 of the Planning Manual, with the trend projections integrated into the baseline through a defined sequence of consistency checking and uncertainty aggregation.

Chapter 17 — Visioning and Backcasting

Normative scenario methods that start from a desired future and work backwards.

17.1 What Visioning and Backcasting Are

Visioning and backcasting are normative scenario methods that begin from a desired future and work backwards to the present. Visioning is the construction of the desired future state itself: the definition of what the planning authority intends the system to look like at the horizon date, expressed in observable conditions rather than in aspiration. Backcasting is the analytical step that follows: the construction of the path from the present to the desired future, the milestones along the path, and the decisions required at each milestone. The two methods are conducted in sequence, with the vision as the input and the backcast as the product that converts the vision into a planning object. Together they constitute the normative branch of foresight, distinct from the exploratory branch treated in the preceding chapters.

The distinction from exploratory scenarios is the distinction between the future the planning authority wants and the futures it might face. Exploratory scenarios, treated in Chapters 9 through 12, start from the present and project forward across the critical uncertainties, producing a range of futures against which the plan is tested for robustness. Normative scenarios start from a chosen future and work backward, producing the path the plan must follow if the future is to be realised. The exploratory method asks what could happen; the normative method asks what must be done. The choice is governed by the planning question, as the Planning Manual's Chapter 7 sets out at the general level and the doctrine develops at the specialty level.

The use in long-term strategic planning arises because some objectives can only be reached by sustained action over horizons that exceed a single plan cycle. Energy transitions, demographic rebalancing, infrastructure systems, and biodiversity recovery operate on horizons of twenty to fifty years, and their attainment depends on decisions taken in successive plans. Visioning defines the objective at the horizon date, and backcasting decomposes the path into segments the planning authority can act on within each cycle. The Natural Step framework, developed in Sweden in the late 1980s, applied this method to sustainability by defining system conditions for a sustainable society and working backwards to the strategic steps required to meet them. The Dutch Sustainable Technology Programme of the 1990s applied the same logic to technology policy, and the European Union's forward-looking studies have applied it across sectoral strategies since.

17.2 The Visioning Method

The vision is constructed from three components: the desired future state, the values it embodies, and the conditions it requires. The desired future state is a description of the system at the horizon date in observable terms: the demographic profile, the composition of the energy system, the structure of the economy, the condition of the natural environment, the capacity of the public services. The values the future embodies are the political and ethical commitments that

distinguish this future from others the planning authority could have chosen: equity, sustainability, security, prosperity, resilience, identity. The conditions the future requires are the structural preconditions without which it cannot be realised: the infrastructure that must exist, the institutions that must operate, the capabilities that must be developed, the resources that must be available.

The participatory character of visioning is structural rather than ornamental. A vision constructed by a planning staff and presented as a finished product lacks the political and social authority that comes from having been built with those who will live in the future it describes. The construction convenes the affected interests, the technical experts, and the political principals, and works through structured engagement to identify the desired future, the values, and the conditions. The Dutch Sustainable Technology Programme's visioning exercises convened industry, research, government, and civil society actors over months to construct shared visions of sustainable technology futures. The practice treats the convening as part of the analytical work, because the participants' contributions shape the vision's content.

The relationship to political direction is constitutive. A vision that does not reflect a political decision about the desired future is a staff aspiration, and a staff aspiration cannot bind a future planning cycle. The vision must be adopted by the political authority that has the power to commit the planning system to the path, recorded in the form the jurisdiction uses for such commitments: a cabinet decision, a legislative resolution, a national development plan, or a sectoral strategy approved at the political level. The Natural Step framework's adoption by Swedish municipalities in the 1990s illustrates the relationship: the framework supplied the analytical structure, and the municipalities' political adoption supplied the authority. The vision is documented with the future state, the values, the conditions, the participants, and the date of adoption, so that it can be reviewed and revised as conditions change.

17.3 The Backcasting Method

Backcasting begins with the documented vision and works backward to the present, constructing the path the planning authority must follow if the vision is to be realised. The path is composed of milestones: observable states of the system, dated and specified, that mark the progress from the present to the horizon. The milestones are set at intervals consistent with the planning cycle, so that each plan adopts the milestones that fall within its term and the subsequent plans adopt those that fall within theirs. At each milestone the backcast identifies the decisions required: the investments to be made, the regulations to be enacted, the institutions to be established, the capabilities to be developed. The backcast's discipline is to identify the decisions explicitly rather than to assume they will be made.

The construction of the path proceeds by backward induction from the vision to the present. The analyst begins at the horizon date and asks what the system must look like five years before, if the vision is to be reached at the horizon; the answer is the next-to-last milestone. The process continues, milestone by milestone, until the path reaches the present, with each step justified by the next rather than by the present. At each step the analyst identifies the decisions required to move from one milestone to the next, recorded in the order the path reverses, so that the resulting schedule runs from the present forward. The backward construction is the method's defining

discipline, because it exposes the dependencies a forward projection would conceal and forces each milestone to be justified by its successor.

The identification of the first steps is the backcast's principal output for the current planning cycle. The first steps are the decisions that must be taken in the current plan if the first milestone is to be reached, and they enter the plan as commitments. They typically include the investments with long lead times, the institutional reforms that require legislation, the capability programmes that take a generation to mature, and the decisions that, if deferred, foreclose future options. The Dutch Sustainable Technology backcasts identified research and development investments, regulatory frameworks, and infrastructure decisions as the first steps, and showed that this discipline was what converted the backcast into a planning instrument. The plan adopts the first steps as commitments and the subsequent milestones as assumptions for the next plan, with both the vision and the backcast revised at intervals consistent with the planning cycle.

17.4 When to Use Normative Methods

Visioning and backcasting are appropriate when three conditions are jointly met: a clear desired future, a long horizon, and a need to align action around a direction. The clear desired future is the political condition: the planning authority must have decided, and recorded, what future it intends to realise, with sufficient specificity that the vision can be backcast. The long horizon is the temporal condition: the future must be far enough from the present that the path cannot be planned in a single cycle and the decomposition into milestones adds analytical value. The need to align action is the organisational condition: the planning authority must need to coordinate the actions of multiple plans, agencies, and cycles toward the same destination, in circumstances where without alignment the actions would diverge.

The conditions are not always met, and the methods are not always appropriate. A planning question without a politically determined desired future cannot be addressed by visioning, because the method requires the future as input; such questions are addressed by exploratory scenarios, which construct the range of plausible futures against which the plan is tested. A planning question with a horizon short enough to be planned in a single cycle does not require backcasting, because the path can be constructed by forward planning without the discipline of backward induction. A planning question requiring alignment across one cycle and one agency does not require the method, because ordinary coordination suffices. The Natural Step framework's application to sustainability planning illustrates the appropriate case: a politically determined future, a horizon of decades, and a need to align action across multiple sectors and jurisdictions.

The methods are also inappropriate when the desired future is fixed in a way that precludes revision, because the vision must be revisable as conditions change. A vision adopted as a permanent commitment would either be abandoned at the first change of government or be enforced against the conditions the future had produced. The doctrine treats the vision as revisable: the political authority that adopts it may revise it, the backcast is updated to reflect the revision, and the plan is adjusted accordingly. The European Union's forward-looking studies on energy and climate illustrate the revisable approach: the 2050 objectives have been adopted, revised, and re-adopted across successive European commissions, and the backcasts updated to track the revisions.

The revisability is a condition of the method's use, because long-horizon planning that cannot learn from changing conditions has abandoned planning for prophecy.

17.5 The Relationship to Exploratory Scenarios

Normative and exploratory methods are complementary rather than alternative. Exploratory scenarios, treated in Chapters 9 through 12, construct the range of plausible futures the plan must be robust against; normative scenarios construct the path to a desired future the plan must follow. The two methods address different questions, but the questions arise together in long-horizon strategic planning. The vision identifies the destination, and the exploratory scenarios identify the conditions the path will pass through. A plan that has a vision but no exploratory scenarios commits to a destination without knowing what the journey will encounter; a plan that has exploratory scenarios but no vision knows what it might encounter but has not decided where it is going.

The principal use of exploratory scenarios in normative planning is to test the vision's robustness: whether the vision can be realised under each of the plausible conditions the future might produce. The test is conducted by backcasting the vision under each scenario in turn, and examining whether the path is feasible under each. A vision whose path is feasible under all the scenarios is robust, and the authority can commit to the path with confidence that no plausible development will foreclose it. A vision whose path is feasible under some scenarios but not others is conditional, and the authority must monitor the conditions and adjust the path as the future unfolds. A vision whose path is infeasible under most scenarios is fragile, and the authority must revise the vision or accept that its attainment depends on a particular development of conditions.

The integration of the two approaches is the mature practice of long-horizon strategic planning. The vision supplies the destination; the exploratory scenarios supply the conditions; the backcast supplies the path; the scenario-based test supplies the assessment of the path's robustness. The integrated product is a strategy that combines a normative commitment to a desired future with an exploratory awareness of the conditions the future may produce. The Dutch Sustainable Technology Programme integrated the two by constructing normative visions of sustainable technology futures and testing the paths to them against exploratory scenarios of technological, economic, and regulatory change. The European Union's 2050 energy and climate strategies integrate the two in the same way, with the vision of a decarbonised energy system tested against scenarios of technological development, fossil-fuel prices, and geopolitical conditions.

17.6 Worked Example: Backcasting from a Renewable Energy Vision

Consider a planning authority that has adopted a vision of a fully renewable electricity system by 2050, with the values of decarbonisation, energy security, and affordability, and the conditions of sufficient generating capacity, adequate transmission, dispatchable backup, and skilled workforce. The horizon is 2050, the present is 2025, and the planning cycle is five years, so the backcast identifies milestones at 2030, 2035, 2040, and 2045, with the first milestone falling within

the current plan. The backcast proceeds by backward induction, beginning at 2050 and working toward the present, with each milestone justified by the next.

At 2050 the system is fully renewable, with the conditions of the vision met. At 2045, five years before, the system must be within reach of these conditions: the last fossil-fuel plants scheduled for closure or conversion, the transmission complete or under final construction, the dispatchable backup in commission, and the workforce training producing the cohorts the system will need. At 2040 the system must be at the point from which the 2045 state can be reached: the fossil-fuel plants not yet closed must be the minority of capacity, the transmission substantially built, the dispatchable backup under construction, and the training programmes at scale. At 2035 the renewable share must be high enough that the 2040 state can be reached within the following cycle, with the investments in transmission and backup underway and the training programmes established.

At 2030, the first milestone, the renewable share must be at the level from which the 2035 state can be reached: renewable generation at scale, transmission planning complete and construction begun, the regulatory framework for dispatchable backup in place, and the workforce training programmes launched. The decisions required between 2025 and 2030 are the first steps the current plan must adopt: the generation tenders, the transmission corridors, the backup regulations, and the training programme funding. The first steps enter the plan as commitments, recorded with the responsible authorities and the dates by which the actions must be taken. The milestones from 2035 onward enter the plan as the assumptions against which the next plan will be framed, with the recognition that their attainment depends on the decisions the current plan has adopted.

The scenario-based test of the vision’s robustness follows. The backcast is conducted under each of the exploratory scenarios for the energy system: rapid technological progress, stagnant progress, high fossil-fuel prices, and low fossil-fuel prices. Under the rapid-progress and high-price scenarios the path is feasible, and the first steps remain as identified. Under the stagnant-progress and low-price scenarios the path is more demanding: the first steps must be expanded to include additional support for research and development and additional protection for the renewable investments against market competition. The test identifies the conditions under which the first steps must be adjusted, and the monitoring system is specified to detect which scenario is emerging so the adjustments can be made in time.

17.7 Chapter Summary

Visioning and backcasting are normative scenario methods that begin from a desired future and work backwards to the present. Visioning is the construction of the desired future state, the values it embodies, and the conditions it requires, conducted through participatory engagement and adopted by political authority. Backcasting is the working backwards from the vision to the present by backward induction, identifying the milestones, the decisions required at each milestone, and the first steps the current plan must adopt. The two methods are conducted in sequence, with the vision as the input and the backcast as the analytical product that converts the vision into a planning

object. They are appropriate when a clear desired future, a long horizon, and a need to align action are jointly present, and inappropriate when any is absent.

The Natural Step framework, the Dutch Sustainable Technology Programme, and the European Union's forward-looking studies on energy, climate, and demographic change constitute the principal comparative references. The Natural Step framework supplies the analytical structure of system conditions and the strategic steps required to meet them, applied to sustainability planning in Sweden and beyond. The Dutch Sustainable Technology Programme supplies the model of participatory visioning and technology-policy backcasting, with industry, research, government, and civil society convened to construct the vision. The European Union's forward-looking studies supply the model of revisable long-horizon objectives tested against exploratory scenarios and updated across successive political cycles. The doctrine treated in the doctrine draws on all three, and the methods are the same across the mature jurisdictions that have applied them.

The complementarity with exploratory scenarios is the mature practice of long-horizon strategic planning. The vision supplies the destination, the exploratory scenarios supply the conditions the path will pass through, the backcast supplies the path, and the scenario-based test supplies the assessment of the path's robustness. The integrated product is a strategy that combines a normative commitment to a desired future with an exploratory awareness of the conditions the future may produce. The worked example of a renewable energy vision illustrates the method's application: the vision is constructed, the backcast identifies the milestones and the first steps, and the scenario test assesses the path's robustness under the plausible conditions the energy system may produce. The chapters that follow treat the integration of normative and exploratory methods in strategic risk assessment and in the design of monitoring systems.

Chapter 18 — Designing and Conducting the Foresight Exercise

Part V: Foresight in Practice — The design and management of foresight exercises.

18.1 What the Foresight Exercise Is

A foresight exercise is the structured application of foresight methods to a defined question, conducted over a defined period and producing a defined foresight product. The exercise is initiated by a commissioning authority that requires foresight input on a specific question, scoped to a specific horizon and scope, and conducted by a designated exercise team under an exercise lead. The product is a foresight output — a scenario set, a trend assessment, a driver map, a strategic risk assessment — that the commissioning authority will use in subsequent strategy or planning decisions. The exercise is bounded by a beginning, an end, and a defined product, and its success is judged by whether the product answers the question posed.

The distinction from continuous scanning is structural. Continuous scanning is the standing function of the foresight unit: the routine collection of signals, the maintenance of the horizon scan log, the periodic update of trend assessments, conducted without a specific question and without a defined end. A foresight exercise is initiated because a question has arisen that the standing function cannot answer at the depth required — a planning cycle is opening, a strategic decision is pending, a disturbance has revealed an unexamined area. The exercise draws on the standing function's accumulated material and returns its findings to the standing function's record, while remaining a discrete undertaking with its own authority, resources, and timeline.

The exercise is conducted as a project, with the discipline that project management imposes. The discipline comprises a defined scope, a defined timeline, defined resources, defined roles, and defined quality standards, all stated in the exercise's terms of reference and maintained by the exercise lead throughout. The project framing ensures the exercise can be held accountable for delivering the product the commissioning authority requires, on the schedule and with the resources allocated. A foresight exercise conducted without project discipline consumes time and expertise without producing a usable product.

Comparative practice confirms the framing. Singapore's Centre for Strategic Futures conducts discrete undertakings, each commissioned, scoped, and delivered against a defined brief. The United Kingdom's Foresight projects, conducted under the Government Office for Science, run for defined periods — typically eighteen months to two years — on a defined topic, with a defined output, and are distinct from the standing horizon scanning function. France's Commissariat Général du Plan, during its operational period, commissioned prospective studies as discrete exercises, each with its own terms of reference and its own product. The pattern across mature jurisdictions is consistent: the exercise is a project, bounded and managed as such.

18.2 The Scoping Phase

The scoping phase is the first phase of the exercise and the most consequential. Scoping defines the question the exercise will address, the horizon over which it will address it, the scope it will cover, the stakeholders it will engage, and the resources it will consume. The output of the scoping phase is the exercise's terms of reference, which record these definitions and serve as the exercise's governing document throughout its conduct. A foresight exercise conducted without a clear scoping produces a product whose relationship to the commissioning authority's question cannot be established.

The question is the scoping's central element, and its framing determines the exercise's entire conduct. A question framed as the future of energy supply produces a different exercise from one framed as the conditions that would require the plan's assumptions about energy supply to be revised. The first invites exploration without end; the second invites analysis directed at a defined planning decision. The exercise lead frames the question so that it is answerable within the horizon and the resources available, bounded enough to permit analysis but open enough to capture the relevant uncertainties. The framing is recorded in the terms of reference, with the reasoning, so the later stages can be tested against the question the exercise was commissioned to answer.

The horizon and the scope are defined together. The horizon is the period over which the exercise will project — typically ten, twenty, or thirty years for a national plan, shorter for a sectoral exercise. The scope is the boundary of what the exercise will consider: the geographic scope, the sectoral scope, the population affected, and the institutional actors involved. A horizon of thirty years with a national scope produces a different exercise from a horizon of ten years with a sectoral scope. The terms of reference record the horizon and the scope, together with the trade-offs — the elements, time periods, and questions the exercise will not address — with the exclusions documented so the commissioning authority knows what the exercise does not cover.

The identification of the exercise lead and the exercise team completes the scoping. The exercise lead is the practitioner responsible for the exercise's conduct and its product, accountable to the commissioning authority, and empowered to make the methodological and resource decisions the exercise requires. The exercise team comprises the practitioners who will conduct the collection, the analysis, and the dissemination, drawn from the foresight function, the relevant sectoral expertise, and the analytical support functions. The terms of reference record the lead's authority and the team's composition, with the roles, time commitments, and reporting lines stated. The team is assembled against the question the exercise will address, and the scoping's quality is judged by whether the team can plausibly produce the required product.

18.3 The Collection Phase

The collection phase gathers the signals, data, and expert input the analysis will require. Collection draws on the methods treated in the preceding chapters — horizon scanning, trend analysis, driver mapping, expert elicitation, structured workshops — applied in the combination the question and the scope require. The phase's output is a body of evidence on the trends, drivers, and uncertainties relevant to the question, organised so that the analysis phase can use it. Collection is the targeted gathering of evidence against the question the terms of reference has posed, conducted

under the exercise lead's management, with every source engaged and every signal recorded required to advance the question.

The management of collection is the exercise lead's principal task during this phase. The lead assigns collection tasks to team members against the question and the timeline, monitors the coverage to ensure the relevant source types are addressed, and adjusts the assignments as the collection reveals what is and is not available. The lead also manages the engagement with external sources — the statistics office, the sectoral ministries, the external experts, the international organisations — to ensure the exercise's requests are coherent and to avoid duplication. Collection that does not advance the question is curtailed, and collection that reveals the question requires refinement is referred back to the commissioning authority for a scoping decision.

The documentation of sources is the discipline that distinguishes a defensible collection from an undirected accumulation. Every signal, data point, expert judgment, and workshop output is recorded with its source, its date, and the analyst's assessment of its reliability and relevance, following the standards treated in Chapter 11 for the horizon scan log and Chapter 14 for expert elicitation. The documentation permits the collection to be audited: a reviewer can trace any finding in the analysis back to the evidence on which it rests, assess whether the evidence is sufficient, and identify the gaps. The documentation also permits the collection to be reused, with the next exercise on a related question beginning from the prior exercise's record rather than from a blank slate.

The phase's conduct proceeds in cycles. The collection reveals gaps that require additional sources; the gaps are addressed through targeted collection, which may reveal further gaps, and the cycles are bounded by the timeline and the resources. The exercise lead's task is to manage the cycles so that the collection converges on a body of evidence sufficient to support the analysis within the constraints. The convergence is judged by whether the collection covers the trends, drivers, and uncertainties the question requires, and by whether the evidence is of sufficient quality to support the analysis. The phase ends when the exercise lead judges the collection sufficient, and the judgment is recorded with its basis so the commissioning authority can assess whether the sufficiency standard has been met.

18.4 The Analysis Phase

The analysis phase converts the collection's evidence into the exercise's analytical product. Analysis identifies the drivers that shape the question's subject, constructs the scenarios that span the plausible futures, and assesses the implications of those futures for the question the exercise was commissioned to address. The phase's output is the draft foresight product — the scenario set, the trend assessment, the strategic risk assessment, or the combination of these the terms of reference has specified. The analysis is the exercise's intellectual core, and its quality is the principal determinant of the exercise's value to the commissioning authority.

The integration of methods is the analysis phase's defining characteristic. A foresight exercise that uses a single method — only trend analysis, or only scenarios, or only expert elicitation — produces a product whose limits are the limits of that method. The mature exercise combines methods: trend analysis establishes the empirical baseline, driver mapping identifies the forces that

will shape the baseline's evolution, cross-impact analysis maps the relationships among the drivers, expert elicitation supplies judgment where evidence runs out, and scenario construction integrates the methods into narratives that span the plausible futures. The methods are selected because they advance the question, and they are combined so that each method's output feeds the next method's input.

The management of analysis is the exercise lead's principal task during this phase. The lead assigns analytical tasks to team members against the question and the timeline, monitors the integration of methods to ensure each method's output is feeding the next, and adjudicates the methodological choices that arise during the analysis. The lead also manages the engagement with the analytical community — the sectoral experts who will review the analysis, the external practitioners who will challenge it, the commissioning authority who will use it — to ensure the analysis is tested before it is finalised. Analysis that does not advance the question is curtailed, and analysis that reveals the question requires refinement is referred back to the commissioning authority for a scoping decision.

The production of the draft foresight product completes the analysis phase. The draft sets out the drivers, the scenarios, the implications, and the assessment the analysis has produced, in the form the terms of reference has specified and at the level of detail the commissioning authority requires. The draft is reviewed within the exercise team for internal consistency, by the sectoral experts for substantive accuracy, and by the commissioning authority for relevance to the question. The review produces comments, which the exercise team addresses through revision, and the revised draft advances to the dissemination phase, where it is communicated to the audiences whose decisions it is intended to inform.

18.5 The Dissemination Phase

The dissemination phase communicates the foresight product to the strategy function and to political authority. Dissemination is the active communication of the product's findings to the audiences whose decisions the product is intended to inform, conducted in the form those audiences can use and at the time those audiences need the input. The exercise lead and the team engage with the audiences to ensure the findings are understood, their implications recognised, and their use in the subsequent decisions secured. The phase's output is the product's uptake: the incorporation of the findings into the strategy or planning decisions the commissioning authority will make. The discipline of dissemination is what ensures the analytical investment made in the prior phases is realised in the decisions the exercise was commissioned to inform.

The adaptation of the product to the audience is the dissemination's central task. The strategy function requires the product in a form that supports its work: the assumptions, the risks, the scenarios, and the implications, presented with sufficient analytical depth to support strategic reasoning. Political authority requires the product in a different form: the central findings, the implications for the jurisdiction, and the choices the findings pose, presented with sufficient brevity and clarity to support political judgment. The same product is communicated differently to the two audiences, with the analytical detail preserved for the strategy function and the policy implications

foregrounded for political authority. The adaptation is conducted without distortion: the findings are the same in both forms, and the differences are in the presentation, not in the substance.

The follow-up completes the dissemination phase and extends into the period after the exercise's formal conclusion. The follow-up tracks whether the product's findings have been incorporated into the strategy or planning decisions, identifies where they have not and why, and feeds back to the foresight function the lessons the exercise has produced. The follow-up also maintains the product for the period the commissioning authority requires: the indicators the analysis identified are monitored, the scenarios are revisited as conditions change, and the product is revised when the conditions warrant. The follow-up is the bridge from the discrete exercise to the standing function, ensuring the exercise's findings remain available to the planning process across the period the exercise addressed.

18.6 Managing the Exercise

The management of the exercise is the exercise lead's continuing responsibility across all four phases. Management comprises the maintenance of the timeline, the deployment of the resources, the conduct of the quality assurance, and the management of the risks, conducted against the terms of reference and reported to the commissioning authority at defined intervals. Management is the frame within which the analytical work is conducted, integral to the exercise and exercised by the lead as a continuing responsibility. The discipline of management ensures the analytical effort produces a usable product on the schedule and with the resources the commissioning authority has allocated.

The timeline and the resources are managed together, because they are the two constraints within which the exercise is conducted. The timeline sets the dates by which the phases must be completed and the product delivered, with the milestones the commissioning authority requires. The resources comprise the staff time, the expertise, the data, the workshop facilities, and the external engagements the exercise requires, allocated across the phases according to the demands each imposes. The exercise lead manages the trade-offs: a question more complex than the scoping anticipated may require an extension of the timeline, a reallocation of resources, or a refinement of the question. Trade-offs within the terms of reference are the lead's to make; those that exceed them are referred to the commissioning authority.

The quality assurance and the risk management are the management's two control functions. Quality assurance is the conduct of review at defined points — the scoping, the collection, the analysis, and the draft product — each reviewed before the next phase begins by qualified reviewers not involved in the work. The findings are addressed before the work advances. Risk management is the identification, assessment, and treatment of the risks the exercise faces: that the question is unanswerable within the resources, that the data is insufficient, that the experts disagree irreconcilably, or that the political context shifts. The risks are identified in the scoping, monitored through the conduct, and treated by the lead or referred to the commissioning authority where the treatment exceeds the lead's authority.

The common pitfalls are documented across mature practice, and their remedies are the disciplines treated in the doctrine. The principal pitfalls are a question framed too broadly, a scope

that excludes critical factors, a collection that exceeds the analysis's capacity, an analysis that produces scenarios without consistency tests, a dissemination that misses the policy implications, and a follow-up omitted when the team disperses. The remedies are a scoping that frames the question against the resources, a collection managed against the question, an analysis that integrates the methods with their consistency tests, a dissemination that adapts the product to the audience, and a follow-up specified in the terms of reference and resourced as a phase of the exercise. The mature exercise lead studies the pattern before beginning a new exercise, so the same errors are not repeated.

18.7 Chapter Summary

The foresight exercise is the structured application of foresight methods to a defined question, conducted as a bounded project with a beginning, an end, and a defined product. The exercise is distinguished from continuous scanning by its initiation against a specific question and by its conduct under defined terms of reference. The exercise is conducted in four phases — scoping, collection, analysis, and dissemination — with a management frame that maintains the timeline, the resources, the quality assurance, and the risk management across all four.

The scoping phase defines the question, the horizon, the scope, the stakeholders, and the resources, and produces the terms of reference that govern the exercise. The collection phase gathers the signals, data, and expert input the analysis requires, managed against the question and documented to a standard that permits the collection to be audited and reused. The analysis phase integrates the methods to produce the draft foresight product, with the integration governed by the question and the management conducted against the terms of reference. The dissemination phase communicates the product to the strategy function and to political authority, adapted to each audience without distortion, with a follow-up that maintains the product's availability across the period the exercise addressed.

The management of the exercise is the exercise lead's continuing responsibility, conducted against the terms of reference and reported to the commissioning authority at defined intervals. The common pitfalls — a question framed too broadly, a scope that excludes critical factors, a collection that exceeds the analysis's capacity, an analysis that omits the consistency tests, a dissemination that misses the policy implications, a follow-up that is omitted — are addressed through the disciplines of scoping, collection management, method integration, audience adaptation, and follow-up specification. The exercise's success is judged by whether the product answers the question the commissioning authority posed, on the schedule and with the resources allocated, and by whether the product's findings are incorporated into the decisions the exercise was commissioned to inform.

Chapter 19 — Foresight Products and Communication

The foresight product and its communication to decision-makers.

19.1 What Foresight Products Are

A foresight product is the document or presentation through which the findings of a foresight exercise are communicated to decision-makers. The product is the output of the dissemination phase of the foresight process, the fourth phase treated in Chapter 6 of the Planning Manual, and it is the only part of the exercise the decision-maker ordinarily sees. The product is distinct from the process: the process is the analytical work of scoping, collection, analysis, and dissemination, conducted over weeks or months; the product is the structured artefact that distils that work into a form the decision-maker can act on. The product carries the foresight function's findings into the decision, and its quality is judged by whether the decision is better informed as a result.

The distinction between product and process is the distinction between communication and analysis. A foresight exercise produces a large body of working material: signal logs, driver maps, cross-impact matrices, scenario narratives, workshop proceedings, and expert elicitations. These materials are the analytical record, and they are maintained as such, but they are not the product. The product is a separate artefact, constructed for an audience that did not participate in the analysis and that will not read the analytical record. The discipline of product construction is the discipline of selection: of choosing, from the analytical record, the findings the decision requires and presenting them in the form the decision-maker can use.

The governing requirement is that the product serves the decision. A foresight product that conveys all of the analysis but does not inform the decision has failed, regardless of its analytical merit; a foresight product that selects rigorously and informs the decision has succeeded, even at the cost of omitting material the analysts found interesting. The requirement is operationalised at the start of product construction, by stating the decision the product is to inform, the audience that will receive it, and the question that audience faces. The product is then constructed to that specification, and the analytical record is held available for the reader who requires deeper detail. The mature practice treats the product as a deliverable with a stated purpose, not as a container for the exercise's outputs.

19.2 Product Types

Five product types are used in mature foresight practice: the foresight report, the scenario set, the horizon scan briefing, the foresight memo, and the emerging issues brief. The foresight report is the comprehensive product of a major exercise, typically thirty to eighty pages, setting out the question, the method, the evidence, the scenarios or trend assessment, the implications, and the recommended handling of uncertainty. The report is the product of record: it is archived, cited, and used as the reference for subsequent exercises. Singapore's Centre for Strategic Futures and the United Kingdom's Government Office for Science both produce reports of this kind, as does

Finland's Government Report on the Future, which is submitted to Parliament under the country's foresight statute. The report's length and structure reflect its purpose as the standing reference for the exercise.

The scenario set and the horizon scan briefing are the principal recurring products. The scenario set is the standalone publication of the scenarios constructed in an exercise, presented with their narratives, their distinguishing assumptions, and the indicators that would signal which scenario is emerging. The scenario set is shorter than the report, typically ten to thirty pages, and is used where the scenarios are the principal output and the supporting analysis is held elsewhere. The horizon scan briefing is the periodic product of a standing scan function, published quarterly or annually, that lists the significant signals detected in the period, assesses their implications, and flags those that warrant escalation. The United Kingdom's horizon scanning programme and the European Commission's European Strategy and Policy Analysis System both produce briefings of this kind, and the form has become standard across mature jurisdictions.

The foresight memo and the emerging issues brief are the focused products. The foresight memo is a short document, typically two to six pages, addressed to a specific decision: it states the foresight finding relevant to that decision, the uncertainty attaching to the finding, and the implications for the decision's options. The memo is the product of choice where a decision is in train and the foresight function has findings that bear on it. The emerging issues brief is a short document, typically three to ten pages, that presents an issue the scan has identified as significant and on which the strategy function has not yet engaged. The brief's purpose is to bring the issue to the strategy function's attention before it becomes urgent, and its distinguishing feature is the lead time it provides.

The choice among the five types is governed by the question and the audience. The report suits a major exercise whose findings will be cited for years; the scenario set suits an exercise whose scenarios are the principal output; the briefing suits the periodic output of a standing scan function. The memo suits a decision in train, and the brief suits an emerging issue the strategy function has not yet engaged. The mature practice selects the type that fits the decision, rather than producing a single house format for all exercises.

19.3 Communicating Uncertainty

The communication of uncertainty is the central craft of foresight product construction. A foresight finding is, by the discipline's definition, an assessment of what could happen, not a statement of what will happen, and the product must convey that character to the reader. False certainty — the presentation of a range as a point, of a scenario as a forecast, or of a judgment as a fact — is the principal failure mode, because it invites the decision-maker to treat the product as a prediction and to commit the plan to a single future that the analysis does not support. The discipline requires that the product state, for each finding, the degree of uncertainty attaching to it and the form that uncertainty takes.

The opposing failure is paralysis. A product that communicates uncertainty without communicating a basis for decision leaves the decision-maker with no course of action, and the foresight function's contribution to the decision is lost. The discipline requires that the product

convey not only the range of futures but also the planner's assessment of how the plan should proceed under that range: which elements are robust across the range, which are scenario-dependent, and what indicators would signal which future is emerging. The product's contribution to the decision is the structured handling of uncertainty, not the announcement of it. A product that conveys uncertainty without structure has discharged the analytical obligation but not the communication obligation, and the reader is left with the problem the foresight function was convened to solve.

Three instruments are used to communicate uncertainty in the product. Ranges are used for quantitative findings, where the analysis produces a distribution and the product reports the central case and the spread; the statistics office's population projections, published with high and low variants, are the model for this practice. Scenarios are used for qualitative findings, where the analysis produces a set of plausible futures and the product reports each with its distinguishing assumptions and implications. Confidence assessments are used for judgments, where the analysis rests on expert elicitation rather than measurement, and the product reports the finding with a stated level of confidence and the basis for that level. The Intergovernmental Panel on Climate Change's calibrated language — virtually certain, very likely, likely, about as likely as not — and the United States National Intelligence Council's confidence gradations are the comparative references for this practice, and the foresight function adopts whichever vocabulary is established in its jurisdiction's analytical community.

19.4 Communicating to Political Authority

The communication of foresight to political authority requires adaptation to the audience. Political authority operates under time constraints, accountability pressures, and the requirement to defend decisions in public, and the product constructed for political authority must accommodate these conditions. The comprehensive foresight report is rarely the appropriate product for political authority; the report is the product of record, and it is read by staff, by analysts, and by the public, but the political decision-maker ordinarily requires a shorter and more focused artefact. The mature practice produces, for political authority, a briefing, a memorandum, or a presentation, drawn from the report and constructed to the decision in hand. The report remains the standing reference; the briefing, the memorandum, and the presentation are the working instruments at the political level.

The briefing is the standing instrument for communication to political authority. A briefing is typically two to four pages, structured around the question the decision-maker faces, the foresight findings that bear on that question, the uncertainty attaching to those findings, and the implications for the decision. The briefing is written for the decision-maker who has limited time and who requires the analyst's judgment, not the analyst's reasoning; the reasoning is held in the analytical record and is available on request. The memorandum is the instrument for a specific decision: it is the foresight memo, adapted to the political audience, and it carries the foresight function's assessment of the decision's options. The presentation is the instrument for a structured engagement: a workshop, a strategy retreat, or a cabinet committee, where the foresight function presents findings and is questioned on them, and where the decision-maker's concerns can be addressed in dialogue.

The management of the political interface is the third element of the discipline. The foresight function's independence from political preference is the condition of its usefulness, and the product must convey the analysis without adjustment to the prevailing political position. The discipline requires that the foresight function present the range of futures the analysis supports, including futures the political authority may find inconvenient, and that the strategy function and the political authority choose among the options the range presents. The interface is managed by the discipline of separation: the foresight function produces the findings; the strategy function frames the options; the political authority decides. A foresight function that adjusts its findings to political preference ceases to perform its function, and the institutional arrangements treated in Chapter 6 of the Planning Manual — the location of foresight at the centre of government, with a reporting line that protects its independence — are designed to maintain that separation.

19.5 The Relationship to the Strategy Function

The relationship between the foresight function and the strategy function is the relationship between the function that produces the range of futures and the function that chooses the objectives and means. The two are sequential — foresight first, strategy second — and the handoff between them is the principal mechanism by which foresight enters the plan. The handoff is the structured transfer of the foresight product from the foresight function to the strategy function, accompanied by the analytical record and by a stated set of implications for the plan's assumptions, risks, and options. The handoff is the formal act that closes the foresight exercise and opens the strategy work, and its documentation is the condition under which the strategy function can be held accountable for its use of the foresight.

The handoff is the formal beginning of the collaboration, not its end. The strategy function, in converting the foresight product into strategic options, will identify questions the foresight product does not answer: gaps in the analysis, scenarios that require further development, or implications that require deeper assessment. The mature practice provides for a return loop, in which the strategy function refers such questions to the foresight function, and the foresight function responds with supplementary analysis or with a further product. The collaboration continues through the strategy formulation, the plan drafting, and the plan's adoption, and it is renewed at each revision of the plan. The relationship is institutional, not transactional, and it is sustained by the proximity of the two functions in the institutional arrangements treated in Chapter 6.

The distinction of roles is maintained through the collaboration. The foresight function's role is to produce the range of futures and the implications; the strategy function's role is to choose the objectives and means within that range. The foresight function does not select the plan's assumptions, set the plan's priorities, or commit the plan's resources; those are the strategy function's responsibilities, and the foresight function that encroaches on them has confused its role. The strategy function does not construct the scenarios, assess the drivers, or set the confidence levels; those are the foresight function's responsibilities, and the strategy function that encroaches on them has compromised the analysis. The documentation of the handoff records which function produced which element, so that the responsibilities are traceable through the planning cycle and the audit trail survives staff turnover.

19.6 Chapter Summary

A foresight product is the document or presentation through which the findings of a foresight exercise are communicated to decision-makers. The product is distinct from the process: the process is the analytical work, and the product is the artefact constructed for an audience that did not participate in the analysis. The governing requirement is that the product serves the decision, operationalised by stating the decision, the audience, and the question at the start of product construction.

Five product types are used in mature practice: the foresight report, the scenario set, the horizon scan briefing, the foresight memo, and the emerging issues brief. The report is the comprehensive product of record; the scenario set publishes the scenarios; the briefing is the periodic scan output; the memo addresses a decision in train; the brief brings an emerging issue to the strategy function's attention. The choice of product is governed by the question and the audience, and the mature practice selects the type that fits the decision rather than producing a single house format for all exercises.

The communication of uncertainty is the central craft of product construction. False certainty is the principal failure mode, and the discipline requires that the product state, for each finding, the degree and form of uncertainty attaching to it. Paralysis is the opposing failure, and the discipline requires that the product convey the structured handling of uncertainty, not merely its announcement. Ranges, scenarios, and confidence assessments are the three instruments, and the comparative references include the statistics office's projection variants, the Intergovernmental Panel on Climate Change's calibrated language, and the United States National Intelligence Council's confidence gradations.

Communication to political authority requires adaptation to the audience, and the standing instruments are the briefing, the memorandum, and the presentation, drawn from the analytical record and constructed to the decision in hand. The management of the political interface rests on the separation of roles: the foresight function produces the findings; the strategy function frames the options; the political authority decides. The relationship between the foresight function and the strategy function is sequential and ongoing, formalised in the handoff and sustained through the planning cycle. The documentation of the handoff records which function produced which element, and is the condition under which the responsibilities are traceable.

Chapter 20 — Institutionalising Foresight: Comparative Practice

How mature jurisdictions institutionalise foresight.

20.1 What Institutionalisation Means

Institutionalisation is the establishment of foresight as a permanent function of government, with dedicated staff, defined methods, and a regular cycle that survives changes of personnel and political leadership. The term distinguishes foresight conducted once, in response to a particular question or a particular administration's interest, from foresight conducted as a standing activity of the state. Institutionalised foresight is characterised by continuity: the methods are documented, the products are archived, the staff carry institutional memory, and the cycle recurs on a schedule that does not depend on the next decision to commission a study. The function acquires, through institutionalisation, the properties that allow it to inform strategy over the long horizons foresight addresses.

The same methods — horizon scanning, scenario construction, driver mapping, cross-impact analysis — are used in ad hoc foresight, and the analytical quality may be equal. What ad hoc foresight lacks is the institutional infrastructure that allows the results to be maintained, updated, and brought to bear on successive planning cycles. A scenario set produced for a single plan and then archived is ad hoc; the same scenario set maintained, revised on schedule, and re-presented to the strategy function at the next cycle is institutionalised. The distinction is the difference between a study and a function, and it is the difference that determines whether foresight informs strategy over the long run.

The requirements for institutionalisation are four. The first is dedicated staff, whose primary assignment is foresight and who carry the institutional memory of past exercises. The second is defined methods: documented protocols for horizon scanning, scenario construction, and the other foresight methods, applied consistently across staff and cycles. The third is a regular cycle of exercises aligned with the planning cycle, and the fourth is a mandate and a budget that establish the function, assign its responsibilities, and provide the resources to discharge them. The absence of any one reduces foresight to a series of projects rather than a function.

20.2 The Permanent Foresight Unit

The permanent foresight unit is the institutional form through which foresight is established as a function. The unit is a defined organisational entity, with a head, a staff establishment, a budget line, and a position in the organisational chart, and its existence does not depend on the convening decision of another authority. The unit carries the function's institutional memory, maintains its methods, conducts its cycle, and produces its products. The form is the same across mature jurisdictions, though the units differ markedly in size, name, and exact placement.

The staffing combines methodological specialists and substantive generalists. The methodological specialists maintain expertise in horizon scanning, scenario construction, cross-impact analysis, and the other foresight methods, and they are the custodians of the protocols and the trainers of new staff. The substantive generalists carry knowledge of the policy domains the unit addresses, and they translate the methods into the questions the planning authority faces. The budget is commensurate with the staffing and the cycle: a unit too small to discharge the cycle has been under-resourced, and a unit too large for its mandate has been allowed to drift. The mandate states what the unit is authorised to do, what it is required to produce, and to whom it reports; the reporting line is the most consequential element, because it determines the unit's access to the strategy function.

The relationship to the planning authority is the relationship to the function the unit serves. The unit's products — horizon scans, scenario sets, driver maps, strategic risk assessments — are inputs to the planning authority's strategy development, and the mandate typically requires the unit to produce them on the schedule the planning cycle requires. The relationship to political authority is more distant. The unit advises the planning authority, which in turn advises political authority; the unit does not advise political authority directly, except where political authority has specifically requested the unit's input. The distance is structural: it preserves the unit's analytical independence from the political cycle, and it preserves political authority's responsibility for the decisions the unit's products inform.

20.3 The Foresight Cycle Aligned with the Planning Cycle

The foresight cycle is the schedule of foresight exercises the unit conducts, and its alignment with the planning cycle is the condition under which the unit's products inform the planning process. The planning cycle has phases — scoping, drafting, consultation, adoption, implementation, review — and each phase has different needs for foresight input. The foresight cycle is designed so that the relevant products are available at the points the planning cycle requires them. The alignment is not automatic: it requires the unit to understand the planning cycle's schedule, and the planning authority to understand the foresight cycle's products.

The timing of foresight exercises is set by the planning cycle's calendar. A horizon scan is conducted in the scoping phase, to identify the conditions the plan must address; a scenario exercise is conducted across the scoping and drafting phases, to test the plan's resilience against the range of plausible futures; a driver map is maintained continuously and refreshed before each cycle, to provide the analytical structure the other exercises draw on; a strategic risk assessment is conducted in the drafting phase, to identify the risks the plan must manage. The cycle is set so that each product is delivered with enough lead time for the planning authority to use it, and the unit's calendar is published in advance so the planning authority knows what to expect and when.

The contribution to each phase of the planning process is specific. In the scoping phase, foresight supplies the assumptions the plan will rest on: the demographic projections, the economic outlook, the technological trends, the geopolitical context. In the drafting phase, foresight supplies the scenario set against which the plan's commitments are tested, and the strategic risk assessment that identifies the conditions under which the plan would fail. In the consultation phase, foresight

supplies the analytical material that informs consultation with sectoral ministries and with the public. In the implementation phase, the indicators developed from the scenarios are monitored, and the monitoring results feed the plan's revision; in the review phase, the foresight cycle begins again, with the horizon scan updated, the scenarios revised, and the driver map refreshed.

The alignment requires institutional discipline on both sides. The unit must deliver its products on schedule, and the planning authority must build the products into its process. Where the unit delivers late, the planning authority proceeds without the foresight input, and the unit's work is filed rather than used. Where the planning authority does not build the foresight input into its process, the unit's work is conducted but has no effect on the plan. The alignment is the condition under which foresight is institutionalised in practice rather than only in form.

20.4 The Integration of Foresight into Strategy

The integration of foresight into strategy is the condition under which the unit's products shape the plan, and it comprises more than the handoff. It is the ongoing collaboration between the foresight function and the strategy function through the planning cycle, with the strategy function drawing on the foresight products at each phase and the foresight function adjusting its products in response. The handoff is the formal transfer of a foresight product from the unit to the strategy function, documented with a note stating what the product is, what question it addresses, what its principal findings are, and what its limitations are. The strategy function acknowledges receipt and assigns responsibility for the product's use. The documentation is what permits the product to be traced from its analytical origin to its use in the plan, and a handoff that is not documented cannot be audited.

The ongoing collaboration is the working relationship between the unit and the strategy function through the cycle. The collaboration extends across the planning process, in a series of structured exchanges through which the strategy function poses questions and the unit responds with analysis, and the unit poses questions and the strategy function responds with the policy context. The collaboration is governed by the same protocols as the rest of the planning process: meetings are minuted, exchanges are documented, and the reasoning is recorded. The collaboration is what allows the foresight products to be tailored to the strategy function's needs, and it is what allows the strategy function to understand the foresight products well enough to use them.

The feedback is the return of information from the strategy function to the unit, on the usefulness of the unit's products and the questions the next cycle should address. The feedback is collected at the end of each planning cycle, as the strategy function reviews the products it received and identifies what was useful, what was not, and what was missing. The unit uses the feedback to revise its methods, its scope, and its cycle, so the next cycle's products are better suited to the strategy function's needs. The feedback loop is the mechanism by which the unit learns, and its absence is the condition under which the unit produces the same products cycle after cycle, whether or not they are used.

The avoidance of the foresight-strategy disconnect depends on the handoff, the collaboration, and the feedback all being conducted as institutional practice. The disconnect is the condition in which the unit produces papers and the strategy function proceeds without them, and it is the

failure mode to which foresight functions are most exposed. The disconnect arises where the handoff is informal, where the collaboration is occasional, and where the feedback is absent. The institutional practices that prevent the disconnect are the practices set out above, and they are the practices by which the integration of foresight into strategy is maintained.

20.5 Comparative Examples

Singapore's Centre for Strategic Futures (CSF), established within the Strategy Group of the Prime Minister's Office, is a frequently cited model of institutionalised foresight. The CSF conducts horizon scanning, scenario construction, and strategic risk assessment, and its products inform the Strategy Group's coordination of the government's long-term planning. The CSF's distinctive contribution to the practice of institutionalisation is the placement: by sitting within the Strategy Group, the unit has direct access to the central coordination function, and its products feed the cross-government planning process without intermediation. The CSF also maintains standing horizon scanning infrastructure and convenes the Strategic Futures Network, a community of foresight practitioners across the Singapore government, illustrating the standing-scan approach treated in Chapter 11.

The United Kingdom's Government Office for Science (GO-Science), headed by the Government Chief Scientific Adviser, is a second model. GO-Science conducts foresight projects on themes chosen for their strategic significance, and its reports are presented to the cabinet and to the relevant departments. The distinctive contribution of GO-Science is the project-based foresight exercise, conducted over a defined period with a defined scope and a defined product, alongside the standing horizon scanning maintained by the office. The GO-Science model demonstrates that institutionalised foresight need not be confined to a single form: the unit can conduct both standing scans and project-based deep dives, and the two forms can coexist within the same institutional structure.

Finland's Committee for the Future, a standing committee of the Parliament, is a third model, distinctive for its placement in the legislature rather than the executive. The Committee assesses the future consequences of policy decisions, conducts horizon scanning on behalf of Parliament, and reports on long-term issues; its existence ensures that foresight products are considered by the legislature as well as the executive, and that the legislature has the analytical capacity to assess the government's long-term plans. France Stratégie, the French government's strategic planning commission, is a fourth model, distinctive for its placement as an independent advisory body reporting to the Prime Minister. France Stratégie conducts foresight on long-term issues, publishes its analyses, and contributes to public debate on strategic questions; its independence from line ministries allows it to address questions that cross departmental boundaries.

The European Commission's Joint Research Centre (JRC), through its foresight activities, and the OECD's International Futures Programme (IFP) are the principal multilateral models. The JRC conducts foresight in support of the European Commission's policy development, and its distinctive contribution is methodological: the JRC has developed and refined foresight methods, including scenario construction and cross-impact analysis, and has disseminated them through training and publication. The OECD's IFP conducted foresight on issues of cross-jurisdictional

significance and provided a forum for foresight practitioners across member states to exchange practice. The multilateral models demonstrate that institutionalised foresight operates at the international level as well as the national, and that the methods and the institutional forms are transferable across jurisdictions.

20.6 What Transfers and What Does Not

The principles of institutionalisation transfer across jurisdictions. The principle that foresight requires dedicated staff, defined methods, a regular cycle, and a mandate and budget is general: it holds in Singapore as in the United Kingdom, in Finland as in France, in the European Commission as in the OECD. The principle that the foresight cycle must align with the planning cycle is general: the planning cycle's phases impose the same timing requirements in any jurisdiction that conducts structured planning. The principle that foresight must be integrated with strategy, through handoff, collaboration, and feedback, is general: the foresight-strategy disconnect is a risk in any jurisdiction, and the institutional practices that prevent it are the same. These principles are the doctrine the doctrine has set out, and they are the principles against which any jurisdiction's practice can be assessed.

The institutional forms do not transfer across jurisdictions. The Centre for Strategic Futures is a product of Singapore's administrative tradition, in which central coordination is concentrated in the Prime Minister's Office and the foresight function is placed where it has direct access to that coordination. The Government Office for Science is a product of the United Kingdom's administrative tradition, in which the Government Chief Scientific Adviser has a defined role and the office's authority derives from that role. The Committee for the Future is a product of Finland's parliamentary tradition, and France Stratégie is a product of the French tradition of independent commissions advising the Prime Minister. Each form is the institutional form its jurisdiction's tradition supports, and is not transferable to a jurisdiction with a different tradition.

The assessment of one's own practice against the principles is the test of institutionalisation. A jurisdiction assessing its foresight function does not ask whether its unit looks like the Centre for Strategic Futures or the Government Office for Science; it asks whether its function has dedicated staff, defined methods, a regular cycle, and a mandate and budget, whether the cycle is aligned with the planning cycle, and whether the function is integrated with strategy through handoff, collaboration, and feedback. Where the assessment finds the principles satisfied, the jurisdiction's foresight is institutionalised, whatever its institutional form. Where the assessment finds the principles not satisfied, the jurisdiction's foresight is not institutionalised, however much its unit resembles a model from another jurisdiction. The principles are the test; the forms are the local solutions, and the local solutions are properly the jurisdiction's to design.

20.7 Chapter Summary

Institutionalisation is the establishment of foresight as a permanent function of government, with dedicated staff, defined methods, a regular cycle, and a mandate and budget. The distinction from ad hoc foresight is structural: ad hoc foresight uses the same methods but lacks the

institutional infrastructure that allows the products to be maintained, updated, and brought to bear on successive planning cycles. The permanent foresight unit is the institutional form, with a staffing that combines methodological specialists and substantive generalists, a budget commensurate with the cycle, a mandate that establishes the unit's authority, and a reporting line that determines the unit's access to the strategy function. The relationship to the planning authority is direct, and the relationship to political authority is mediated by the planning authority.

The foresight cycle is aligned with the planning cycle, so that the unit's products are available at the points the planning cycle requires them. The timing of the exercises is set by the planning cycle's calendar, and the contribution to each phase is specific. The integration of foresight into strategy is the condition under which the unit's products shape the plan, and it is constituted by the handoff, the ongoing collaboration, and the feedback between the foresight function and the strategy function. The avoidance of the foresight-strategy disconnect depends on these three practices being conducted as institutional practice.

The comparative examples — Singapore's Centre for Strategic Futures, the United Kingdom's Government Office for Science, Finland's Committee for the Future, France Stratégie, the European Commission's Joint Research Centre, and the OECD's International Futures Programme — illustrate the range of institutional forms the principles support. The principles of institutionalisation, alignment with the planning cycle, and integration with strategy transfer across jurisdictions; the institutional forms do not, because each is a product of its jurisdiction's administrative tradition. The assessment of one's own practice is conducted against the principles, not against the forms, and the principles are the doctrine the doctrine has set out.

Appendix A — The Horizon Scan Log Template

This appendix provides the template for the horizon scan log, the foresight function's institutional memory. The template states the structure, the fields, and the maintenance protocol.

A.1 The Log's Purpose

The horizon scan log is the durable record of the scanning function, maintained as the foresight function's institutional memory. Each entry records a signal detected, assessed, and held for use, and the log as a whole records the scan's accumulated work over time. The log's purpose is to preserve the scan's products across personnel turnover, planning cycles, and changes of government, so that the work of past scanners remains available to present and future users. The log's structure and maintenance are governed by the protocol set out in this appendix.

The log's role in continuity across personnel is its principal institutional function. Scanners transfer, retire, and join; the log remains, and the work accumulated by predecessors is not lost with the staff who produced it. A new scanner joining the team reviews the log to understand what has been scanned, what has been assessed, and what has been escalated, and builds on the work of predecessors instead of starting afresh. The foresight bodies that have sustained scanning over decades — Singapore's Centre for Strategic Futures, the United Kingdom's Government Office for Science, and the secretariats supporting Finland's Committee for the Future — have done so in part because their scan records persisted as staff turned over.

The log's contribution to the analysis phase is the third component of its purpose. The foresight function's analytical exercises — driver mapping, critical uncertainty analysis, scenario construction, and strategic risk assessment — draw on the log for the signals on which they rest. Each exercise consults the log to identify signals relevant to its question, to assemble the evidence for the driving forces it will assess, and to trace the development of those forces over time. The log's structure permits the analyst to find signals by category, by horizon, by source, and by emerging issue, and to follow the network of links among related entries.

A.2 The Log's Structure

The log's structure is defined by the fields each entry contains and by the relationships among those fields. An entry comprises nine fields: date, source, signal description, category, potential impact, uncertainty, time horizon, assessment narrative, and disposition. The fields are fixed across the log, applied to every signal assessed, so that every entry is comparable to every other and the log forms a uniform record rather than a collection of notes. The fields are the minimum content required for an entry to function as institutional memory, and an entry that lacks any field is held for completion before it is filed.

The relationships among the fields constitute the structure within the structure. The date and source fields establish the entry's provenance; the signal description and assessment narrative fields carry its substantive content; the category, potential impact, uncertainty, and time horizon fields

carry its classification; and the disposition field carries its status in the scan's workflow. The classification fields are populated by the scanner's assessment, recorded in the narrative, and the narrative is what makes the classification interpretable by anyone who was not present at the assessment. The disposition field links the entry to the workflow — signals are held, monitored, investigated, escalated, archived, or reactivated — and records the entry's current state in that workflow.

The format of the log is the medium in which the fields and their relationships are maintained. The mature practice sustains the log as a structured database, with fields defined as columns, entries as rows, and links among entries maintained as relations. A spreadsheet is the simpler form, suitable for a smaller scan with a single team and a moderate volume of entries, and a database is the more capable form, suitable for a distributed scan with multiple teams and the need for complex queries. A structured document — a set of files organised by category or by emerging issue — is the simplest form, suitable for a scan being established or operating at small scale, and the form is transitional as the scan grows. The choice of format is governed by the scan's scale, the team's technical capacity, and the institution's information management standards, and the format is reviewed as the scan develops.

A.3 The Fields Defined

The fields are defined to a common standard, applied to every entry in the log. The standard for each field states what the field contains, the form in which it is recorded, and the quality required for the entry to be complete. The standards are documented in the scan's protocols and applied consistently across scanners, so that the log's entries are comparable and the log forms a uniform record. An entry that does not meet the standard for any field is returned to the scanner for completion, and the team leader reviews a sample of entries to confirm consistency.

The date field records the date of the signal's entry into the log, in ISO 8601 format (YYYY-MM-DD), and is the field by which the log is searched for signals detected within a given period. The source field records the publication, dataset, expert, or institution from which the signal was detected, with bibliographic detail sufficient for the source to be located and assessed for quality and bias. The signal description field records the signal in a single sentence — what was observed and where — and serves as the entry's title in list view. The category field records the STEEPL classification — social, technological, economic, environmental, political, or legal — with secondary categories noted where the signal crosses sectors, and the potential impact, uncertainty, and time horizon fields record the scanner's classification of the signal against documented criteria: impact as high, medium, or low; uncertainty as high, medium, or low against criteria distinguishing sparse data, conflicting sources, and genuine indeterminacy; and time horizon as near-term (within the plan horizon), medium-term (within ten years), or long-term (beyond ten years). The assessment narrative field records the scanner's written account of the signal, its source, its implications, and the reasoning behind the classification, in a paragraph or two written to be read by someone who is not the scanner, and the disposition field records the entry's status in the scan's workflow — held, monitored, investigated, escalated, archived, or reactivated — with the date of the last change and the basis for it.

An example illustrates the standard. The date is recorded as 2025-03-14; the source is a peer-reviewed article in a demography journal, with full bibliographic detail; and the signal description reads, “Fertility rate in the jurisdictions covered by the plan fell below 1.3 in 2024, the lowest recorded level.” The category is social, with economic as a secondary category; the potential impact is high, because the implications for labour supply, pension systems, and service demand are substantial; and the uncertainty is medium, because the data are reliable but the trajectory’s persistence is unclear. The time horizon is long-term, because the implications arise beyond ten years, and the assessment narrative sets out the data, the source’s quality and bias, the scanner’s assessment of the signal’s reliability, and the scanner’s judgment about the signal’s significance for the plan. The disposition is investigated, with the date of the last review and the basis for the disposition recorded, and a reader who was not present at the assessment can interpret the entry and act on it.

A.4 The Maintenance Protocol

The maintenance protocol governs the entry schedule, the review schedule, the archival protocol, and the reactivation protocol, and assigns responsibility for each. The entry schedule is the cadence at which new signals are entered into the log: signals are entered as they are detected, in the regular scanning sessions the team conducts, and the entry is completed within the session or, where verification is required, within a defined period thereafter. The entry schedule ensures the log reflects the scan’s current work and that signals do not accumulate in scanners’ notes without entering the shared record. The team leader monitors the entry rate to confirm that the scan is being conducted at the cadence the protocol requires.

The review schedule is the cadence at which the log’s entries are reassessed. Active signals are reviewed monthly, with the scanner confirming whether new information has arrived and whether the signal’s classification or disposition requires revision. The full log is reviewed quarterly, with the team assessing the scan’s coverage, the distribution of entries across categories and horizons, and the state of the prioritisation matrix. The scan’s scope and source list are reviewed annually, timed to feed the planning cycle’s foresight input so that the scan’s coverage is reassessed as each plan period begins.

The archival protocol governs the removal of resolved signals from the active log, and the reactivation protocol governs their return when conditions change. A signal that has been resolved — the development has occurred, the implications have been addressed in the plan, or the signal has been conclusively contradicted — is archived, with the basis for the resolution recorded in the entry. The archival removes the signal from the active log while preserving it in the record, so that the active log remains focused on signals that warrant continued attention. An archived signal that becomes relevant again — because a new development makes it pertinent, because a new question arises, or because a related signal appears — is reactivated, with its previous assessment and history preserved, and brought back into the active log for renewed attention. The team leader is responsible for the schedule, the standards, and the documentation, and the scanner who detected the signal is responsible for the entry; the team leader reviews the team’s work and is accountable to the foresight unit’s head for the log’s integrity.

A.5 Chapter Summary

The horizon scan log is the durable record of the scanning function and the foresight function's institutional memory. The log's purpose is to preserve the scan's accumulated work across personnel turnover, planning cycles, and changes of government, and to provide the source from which the foresight function's analytical exercises draw their signals. Driver mapping, critical uncertainty analysis, scenario construction, and strategic risk assessment each consult the log for the signals on which they rest, and the log's structure permits the analyst to find and trace those signals across the record.

The log's structure is defined by nine fields — date, source, signal description, category, potential impact, uncertainty, time horizon, assessment narrative, and disposition — and by the relationships among them. The fields are fixed across the log, applied to every entry, so that the log forms a uniform record. The fields are defined to a common standard, documented in the scan's protocols and applied consistently across scanners. The format is the medium in which the fields are maintained: a database for a distributed scan at scale, a spreadsheet for a smaller scan, and a structured document for a scan being established.

The maintenance protocol governs the entry schedule, the review schedule, the archival protocol, and the reactivation protocol, and assigns responsibility for each. Signals are entered as they are detected; active signals are reviewed monthly, the full log quarterly, and the scope and source list annually. Resolved signals are archived with the basis recorded, and archived signals are reactivated when conditions change. The team leader is responsible for the schedule, the standards, and the documentation, and the scanner who detected the signal is responsible for the entry. The maintenance protocol is the condition under which the log functions as institutional memory, and a log that is maintained remains a living record.

Appendix B — The Scenario Construction Worksheet

This appendix provides the worksheet for constructing scenario narratives from the critical uncertainty matrix. The worksheet guides the practitioner through the development of each scenario in the set.

B.1 The Worksheet's Purpose

The scenario construction worksheet is a structured form on which the analyst records the components of each scenario in the set as that scenario is developed. The worksheet's function is to ensure that every scenario is constructed against the same fields and to the same standard, so that the set as a whole is comparable across its members. The worksheet is the working record of construction: it captures the analyst's reasoning at the time each scenario is built, and it preserves that reasoning for review, revision, and transmission to the planning authority. The worksheet complements the scenario narrative treated in Chapter 12, by supplying the structured framework against which the narrative is drafted and against which the consistency of the narrative is tested.

The worksheet's role in ensuring consistency across the set is its principal contribution to construction. Consistency, treated in section 12.3, requires that each scenario respect the predetermined elements, the driver relationships recorded in the driver map, and the structure of the critical uncertainty matrix. The worksheet enforces this consistency by requiring the analyst to record, for each scenario, the resolution of the critical uncertainties, the predetermined elements that apply, and the key features that distinguish the scenario from the others in the set. A scenario recorded on the worksheet is exposed to comparison with the others, and inconsistencies among scenarios are easier to detect in the structured fields than they are in narrative form alone.

Documenting the construction process is the worksheet's second principal contribution, and it bears directly on the maintenance of the set over time. A scenario set that is maintained across cycles, treated in section 12.5, requires that the construction record be available to those who revise the set. The worksheet supplies that record: the fields on which the scenario was constructed, the reasoning that supported the construction, and the consistency tests that the scenario was required to pass. When the set is revised, the worksheets for the existing scenarios are reviewed, and the revisions are documented against the original construction. The worksheet is therefore the documentary foundation of the set's maintenance, and its retention is the condition under which the set can be carried forward across staff turnover and exercise cycles.

B.2 The Worksheet's Structure

The worksheet comprises six fields, each of which records a distinct component of the scenario. The fields are: the scenario name; the critical uncertainty resolution, recording the outcomes of the critical uncertainties that define the scenario's cell in the matrix; the narrative description, recording the scenario narrative as treated in section 12.2; the key features, recording the distinguishing characteristics of the scenario; the implications, recording the consequences of

the scenario for the planning question; and the indicators, recording the observable measures that would signal the scenario's emergence. The six fields correspond to the components of construction treated in Chapter 12 and to the use of scenarios in monitoring treated in Chapter 13. Each field performs a distinct function, and the omission of any weakens the worksheet's capacity to support both construction and use.

The fields are arranged in a sequence that reflects the order of construction. The name and the critical uncertainty resolution are recorded first, because they define the cell of the matrix from which the scenario is built. The narrative description is recorded next, because it develops the cell into a coherent future. The key features and the implications are recorded after the narrative, because they are derived from the narrative and depend on it. The indicators are recorded last, because they are derived from the implications and from the critical uncertainty resolution: the indicators are the observable measures that distinguish the scenario from the others, and they are specified only after the scenario's distinguishing features have been settled.

The relationships among the fields are governed by the principle that each field is derived from those that precede it and constrains those that follow. The critical uncertainty resolution constrains the narrative: the narrative must describe a future in which the resolution holds. The narrative constrains the key features and the implications: the features are the salient points of the narrative, and the implications are the consequences the narrative entails. The implications, together with the critical uncertainty resolution, constrain the indicators: the indicators are specified to detect the conditions the implications describe. A change in one field requires a review of the fields that depend on it, and the worksheet's structure is designed to make these dependencies visible to the analyst and to the reviewer.

B.3 The Fields Defined

The scenario name is a short phrase that condenses the scenario's character. The name's function is identification: it allows the scenario to be referred to in the planning process, in the resilience testing, and in the monitoring reports, without restating the narrative. The standard for entry is that the name be distinctive within the set, that it be neutral in connotation, and that it not presuppose the desirability of the future it names. The quality requirement is that the name be memorable and accurate: a name that is memorable but inaccurate misleads, and a name that is accurate but forgettable fails its identifying function. The names used in the worked example of section 12.6 — "Reskilled Transition", "Disrupted Transition", "Prepared Stability", "Drift" — illustrate the standard.

The critical uncertainty resolution records, for each critical uncertainty in the matrix, the outcome that defines the scenario's cell. The resolution's function is to anchor the scenario in the matrix and to specify the conditions under which the scenario differs from the others in the set. The standard for entry is that every critical uncertainty be resolved, that the resolution be stated in the terms used in the matrix, and that the resolution be unambiguous. The narrative description is the scenario narrative, comprising the opening, the trajectory, the end-state, and the implications as treated in section 12.2. The standard for the narrative is internal consistency, vividness, and

accessibility; the quality requirement is that the narrative account for the drivers in the driver map, respect the predetermined elements, and remain consistent with the relationships among drivers.

The key features are the salient points of the scenario, recorded as a short sequence of statements that a reader can grasp without reading the full narrative. The features' function is to support the comparison of scenarios across the set and to support the resilience testing of plans. The standard for entry is that the features be derived from the narrative, that each feature be stated in a single sentence, and that the features collectively distinguish the scenario from the others. The implications record the consequences of the scenario for the planning question, and the indicators record the observable measures that would signal the scenario's emergence, as treated in section 13.3. The quality requirement for the implications is that they bear on the planning question; the quality requirement for the indicators is that they discriminate among scenarios, that they be measurable, and that each be assigned a source, a threshold, and a reporting line.

B.4 The Construction Process

The construction process begins with the critical uncertainty matrix and the driver map, both treated in Chapter 11. The analyst takes the matrix, identifies the cells to be developed into scenarios, and opens a worksheet for each cell. The first two entries on each worksheet are the critical uncertainty resolution, drawn directly from the cell, and the scenario name, chosen to condense the cell's character. With the matrix and the driver map before the analyst, the narrative description is drafted, working from the cell's resolution, the predetermined elements common to all scenarios, and the relationships among drivers recorded in the map. The drafting is iterative: the analyst writes, tests against the map, revises, and tests again, until the narrative is internally consistent.

The worksheet is then completed for the remaining fields. The key features are extracted from the narrative, with each feature stated in a single sentence and the set of features reviewed for distinctiveness within the set. The implications are drawn from the narrative and from the planning question that the scenario set is built to inform; the implications are stated in terms that bear on the question. The indicators are specified, drawing on the critical uncertainty resolution and the implications: for each uncertainty, the indicator is the observable measure whose movement would signal that the corresponding outcome is emerging. The indicators are recorded with their source, threshold, and reporting line, in the form required by the monitoring system treated in Chapter 20 of the Planning Manual.

The completed worksheets are reviewed as a set. The review tests the coverage of the matrix, the symmetry of the set, and the consistency of the narratives with the driver map and with one another. The review also tests the indicators for their power to discriminate among scenarios: an indicator that does not discriminate is replaced. The review is conducted by the construction team and, where the set is to inform a significant decision, by a reviewer outside the team. The review produces a record of revisions, and the worksheets are revised accordingly; the revised worksheets are the documentary product of construction, retained as the construction record and transmitted to the planning authority alongside the scenario narratives.

B.5 Chapter Summary

The scenario construction worksheet is the structured form on which the components of each scenario are recorded as the scenario is built. The worksheet's purpose is to ensure that every scenario in the set is constructed against the same fields and to the same standard, and to document the construction process for review, revision, and transmission. The worksheet complements the scenario narrative by supplying the structured framework against which the narrative is drafted and tested. Its contribution to consistency across the set and to the documentation of the construction process is the principal justification for its use.

Six fields compose the worksheet: the scenario name, the critical uncertainty resolution, the narrative description, the key features, the implications, and the indicators. The fields are arranged in the order of construction, and each field is derived from those that precede it and constrains those that follow. The standards for entry and the quality requirements for each field are specified, and the mature practice applies them consistently across the set. The construction process moves from the matrix and the driver map through the drafting of the narrative to the completion of the remaining fields, with iteration and review at each stage.

Review of the completed worksheets as a set tests coverage of the matrix, symmetry of the set, internal consistency, and the discriminating power of the indicators. The revised worksheets are the documentary product of construction, retained as the construction record and transmitted to the planning authority alongside the narratives. The worksheet's structure and process extend the general doctrine of scenario methods treated in Chapter 7 of the Planning Manual and the specialty craft of construction treated in Chapter 12 of the doctrine. The worksheet is the working instrument by which that craft is brought to bear on the construction of each scenario set.

Appendix C — Expert Elicitation Protocols

This appendix provides protocols for the three principal expert elicitation methods: the Delphi method, the expert panel, and the structured workshop.

C.1 The Delphi Protocol

The Delphi method is a structured elicitation procedure in which a panel of experts answers a written questionnaire over two or more rounds, with controlled feedback after each round, until responses converge or a defined stopping rule is met. The procedure is conducted under anonymity: each expert responds without knowing the identities of the other panellists and without seeing individual responses, and the panel's distribution is the only feedback. The instrument is a questionnaire designed to elicit either quantitative estimates (probabilities, dates, magnitudes) or qualitative judgements (drivers, scenarios, priorities) on a defined question. The questionnaire design is governed by three standards: precision, so that two experts reading the same question understand it the same way, with terms defined and units specified; pre-testing, with pilot respondents drawn from a population similar to the expert panel; and stability, so that revisions across rounds probe the sources of disagreement rather than introduce new questions.

Expert selection is governed by three standards. The panel must be large enough to produce a stable distribution, with twelve to twenty panellists the typical range for a policy Delphi, and diverse across the disciplines, institutions, and perspectives relevant to the question. Each expert must be qualified on the question at hand, with the qualification documented and retained in the elicitation record. The rounds are governed by the standards of controlled feedback and a defined stopping rule: after each round, the panel is provided with the median, the interquartile range, and a summary of the reasoning offered, and is invited to revise; the experts whose responses fall outside the interquartile range are invited to state their reasoning, which is fed back anonymously in the next round.

The stopping rule is set before the first round: the exercise ends when the interquartile range narrows below a threshold defined in advance, when the change between rounds falls below a threshold, or when a fixed number of rounds has been completed. Three rounds are common; more than four rarely produce further convergence. The convergence is reported as a distribution, not as a single number: the median and the interquartile range state the central tendency and the residual disagreement, and the reporting of a Delphi result without the dispersion is a reporting failure. The documentation standards require that the questionnaire, the responses, the summary statistics, the revisions, and the final distribution are retained in the elicitation record, with the reasoning that supports the final positions.

C.2 The Expert Panel Protocol

The expert panel is a structured elicitation procedure in which a defined group of experts meets, under the direction of a facilitator, to address a defined question and to produce a collective

judgement. The panel is distinguished from the Delphi method by the direct interaction among the experts: each expert hears the others and may challenge, support, or qualify the others' positions, and the collective judgement is formed through the interaction. The panel is distinguished from an unstructured committee by the protocol that governs its conduct: the question is defined in advance, the agenda is structured, the facilitator directs the discussion, and the output is recorded. The panel composition is governed by three standards: relevance, so that each member is qualified on the question; diversity, so that the relevant disciplines, institutions, and perspectives are represented; and balance, so that no single school or interest dominates. The panel size is typically six to twelve members: small enough for genuine discussion, large enough to carry the range of expertise the question requires.

The briefing materials are governed by the standards of completeness, balance, and timeliness. Each panel member receives the same briefing pack, comprising the question, the background evidence, the relevant data, and the relevant analyses, with sources cited and limitations stated; the pack is balanced, presenting the available evidence without pre-empting the panel's judgement; and the pack is delivered sufficiently in advance that members can read it before the meeting, typically two weeks. The facilitation is governed by the standards of structure, neutrality, and discipline. The facilitator holds the agenda, ensures that each member speaks to the question, prevents any member from dominating, and records the points made and the points of disagreement; the facilitator is neutral on the substance and does not advance a position.

The recording is governed by the standard of completeness: the proceedings are recorded in full, by minute or by audio, with the speakers identified and the points attributed; the recording is the basis for the synthesis and the audit. The synthesis is governed by the standard of faithfulness: the panel's collective judgement is stated as the synthesis of the points the members made, with the points of agreement and the points of disagreement recorded; where the panel's collective judgement is not unanimous, the majority position and the minority position are both stated, with the reasoning that supports each. The synthesis is presented to the panel for confirmation before it is finalised, and the panel members are given the opportunity to correct the record and to record additional views. The documentation standards require that the briefing pack, the membership, the agenda, the recording, and the synthesis are retained in the elicitation record.

C.3 The Structured Workshop Protocol

The structured workshop is a facilitated elicitation procedure in which a defined group of participants works through a structured agenda to address a defined question, producing outputs that are recorded and synthesised. The workshop is distinguished from the expert panel by the breadth of participation and the activity-based format: the participants include experts and stakeholders, and the agenda comprises structured exercises — scenario construction, driver mapping, cross-impact analysis, priority-setting — rather than open discussion. The workshop is distinguished from an unstructured meeting by the protocol that governs its conduct: the design is set in advance, the participants are selected for their contribution to the question, the facilitator directs the exercises, and the outputs are documented as they are produced. The design is governed by three standards: alignment, so that each exercise addresses a defined sub-question and the

sequence of exercises addresses the workshop's question; realism, so that the exercises can be completed within the time available; and tractability, so that the outputs of each exercise are recorded in a form that supports synthesis.

Participant selection is governed by three standards. The participants must be relevant to the question, with the experts providing the substantive knowledge and the stakeholders providing the perspective of those affected; the participants must be diverse across the disciplines, institutions, and interests the question implicates; and the participant count must match the workshop design, with smaller groups (eight to fifteen) suited to depth and larger groups (twenty to forty) suited to breadth. The facilitation is governed by the standards of structure, neutrality, and pace. The facilitator holds the agenda, directs the exercises, manages the time, and ensures that each participant contributes and that no participant dominates; the facilitator is neutral on the substance and does not advance a position. Where the workshop divides into small groups, each group has a facilitator who applies the same standards within the group, and the small groups report back to the plenary.

The documentation is governed by the standards of completeness and contemporaneity. The outputs of each exercise are recorded as they are produced, on flipcharts, in templates, or in a shared digital workspace, with the recording visible to the participants so they can verify it; the recording is supplemented by a minute of the proceedings, noting the participants, the exercises conducted, the points made in plenary, and the decisions taken. The synthesis is conducted after the workshop, by the workshop team, and is governed by the standard of faithfulness: the synthesis states the outputs of the exercises, the points of agreement and the points of disagreement, and the conclusions the workshop reached, with the reasoning that supports each. The synthesis is presented to the participants for comment before it is finalised, and the comments are incorporated or recorded. The documentation standards require that the design, the participant list, the briefing materials, the exercise outputs, the minute, and the synthesis are retained in the elicitation record.

C.4 The Elicitation Record

The elicitation record is the documentation retained from an expert elicitation, comprising the method, the participants, the inputs, the outputs, and the synthesis, sufficient for a qualified reader to understand what was done, by whom, with what evidence, and with what result. The record serves two purposes: transparency and institutional memory. The transparency purpose is to permit the elicitation's results to be scrutinised, by the planning authority that commissioned the elicitation, by the auditors who examine the planning process, and by the public to whom the plan is ultimately accountable; the scrutiny depends on the record's completeness, because an elicitation whose record is incomplete cannot be assessed against the standards. The institutional memory purpose is to permit the elicitation's results to be revisited and revised in subsequent planning cycles, by a team that may include none of the original members; the revisiting depends on the record's completeness for the same reason.

The record's contents are governed by the standard of completeness across the five elements. The method element states the elicitation procedure used, the protocol applied, and any departures from the standard protocol, with the reasons for the departures. The participants element states the

names, affiliations, and qualifications of the experts or participants, the basis on which they were selected, and any conflicts of interest declared. The inputs element states the questions posed, the briefing materials provided, and the data and analyses the participants were given; the outputs element states the responses, the reasoning offered, and the points of agreement and disagreement; the synthesis element states the collective judgement reached and the basis on which it was reached. Each element is documented to a level of detail that permits a qualified reader to reconstruct the elicitation and to assess its conduct against the protocol.

The record's retention is governed by the standards of duration, accessibility, and integrity. The record is retained for the period the planning authority's archive rules require, which is typically the life of the plan the elicitation informed and one further planning cycle, so the record is available for the plan's review and for the design of the next cycle's elicitation. The record is accessible to those authorised to consult it, with the participants' identities and individual responses protected where the protocol promised anonymity, and the record is stored in a form that preserves its integrity, with version control on the synthesis and timestamps on the inputs. The record's use is governed by the standard of attribution: any quotation from the record in subsequent documents is attributed to the elicitation, the date, and the method, and the elicitation's results are not presented as the planning authority's own judgement without that attribution. The record is the institutional asset that makes the elicitation durable, and its absence is the condition under which the elicitation's results are unusable beyond the cycle in which they were produced.

C.5 Chapter Summary

The three principal expert elicitation methods each have a defined protocol. The Delphi method elicits judgements through iterative rounds of a written questionnaire, conducted under anonymity, with controlled feedback and a stopping rule, and reports its results as a distribution. The expert panel elicits judgements through direct interaction among a defined group, governed by a facilitator, and reports its results as a synthesis of the points of agreement and disagreement. The structured workshop elicits judgements through activity-based exercises conducted by a group of experts and stakeholders, governed by a facilitator, and reports its results as a synthesis of the exercise outputs.

Each protocol is governed by standards that distinguish a defensible elicitation from an aggregation of uninstructed opinions. The Delphi standards address questionnaire precision, panel diversity and size, controlled feedback, and a stopping rule set in advance. The panel standards address panel composition, briefing materials, facilitation, recording, and synthesis. The workshop standards address design, participant selection, facilitation, and documentation. The standards are the basis on which an elicitation is assessed, and the assessment is conducted against the standards, not against the elicitation's apparent sophistication.

The elicitation record is the documentation that makes the elicitation transparent and durable. The record comprises the method, the participants, the inputs, the outputs, and the synthesis, documented to a level of detail that permits reconstruction and assessment. The record serves the transparency purpose, by permitting scrutiny of the elicitation's conduct, and the institutional memory purpose, by permitting the elicitation's results to be revisited in subsequent cycles. The

record is retained for the period the planning authority’s archive rules require, accessible to those authorised to consult it, and stored in a form that preserves its integrity. The record is the institutional asset that makes the elicitation durable, and its absence is the condition under which the elicitation’s results are unusable beyond the cycle in which they were produced.

Appendix D — Foresight Product Templates

This appendix provides templates for the principal foresight products: the foresight report, the scenario briefing, the horizon scan summary, and the emerging issues brief.

D.1 The Foresight Report Template

The foresight report is the comprehensive document that records a foresight exercise's design, conduct, findings, and recommendations. The template comprises six standard sections: the executive summary, the methodology, the findings, the scenarios, the implications, and the recommendations. Each section has defined content and standards, and the report's authority rests on the consistency with which the sections are populated across exercises. The template is published, dated, and version-controlled, so that reports produced at different times and by different teams remain comparable to one another and to reports produced elsewhere.

The executive summary states the question examined, the horizon, the range of futures identified, the principal implications for the plan, and the recommended actions, in fewer than 1,200 words. The methodology section documents the sources scanned, the experts consulted, the methods applied, the workshops conducted, and the assumptions made in selecting and combining inputs. The findings section presents the evidence base, organised by driver, with each driver assessed for trajectory, uncertainty, and relevance to the question. The scenarios section presents the constructed futures in the format specified in the scenario briefing template, with the indicators that distinguish them.

The implications section translates the scenarios into consequences for the plan's assumptions, risks, and revision triggers, with each implication traceable to the scenario from which it is drawn. The recommendations section sets out the actions the planning authority should take in response, distinguishing actions to be taken now from actions contingent on the emergence of specific indicators. Each section is written so that it can be read on its own and so that a reader can navigate between sections by reference. The report is reviewed by a panel that includes at least one practitioner not involved in the exercise, and the review record is appended to the published report.

D.2 The Scenario Briefing Template

The scenario briefing is the short document that presents a constructed scenario set to decision-makers, distinct from the foresight report in length and purpose. The template comprises four standard sections: the scenario set, the key features of each scenario, the indicators, and the implications for the plan. The briefing is typically eight to twelve pages, written for a non-specialist audience, and updated on a schedule determined by the monitoring cycle. The format is fixed so that successive briefings are directly comparable and so that decision-makers can locate the information they need without re-learning the structure.

The scenario set section states the question examined, the horizon, the axes of critical uncertainty used to construct the set, and the number of scenarios produced, with a one-paragraph

rationale for the set design. The key features section presents each scenario in a single page, with a short narrative, a statement of the conditions that define the scenario, and a note on the scenario's plausibility relative to the present. The indicators section lists, for each scenario, the early signals that would mark its emergence, with a threshold specified for each indicator at which the scenario is judged to be materialising. The implications section translates the scenarios into consequences for the plan, identifying which assumptions are at risk under each scenario and which decisions would need to be revised.

The briefing's standard is that each scenario must be plausible on its own terms, internally consistent across its features, and distinct from the others in the set. Plausibility requires that the scenario's conditions can be defended against the evidence, consistency requires that the features do not contradict one another, and distinctness requires that the scenarios do not collapse into variants of a central case. The briefing carries a date, a version number, and a statement of the next scheduled update. Decision-makers receive the briefing alongside the monitoring report that tests the indicators against current conditions.

D.3 The Horizon Scan Summary Template

The horizon scan summary is the periodic document that records the signals identified during a scan cycle and the emerging issues derived from them. The template comprises four standard sections: the period covered, the signals identified, the emerging issues, and the assessment. The summary is produced on a fixed schedule, typically quarterly, and is shorter than the foresight report, typically ten to twenty pages. The format is fixed so that successive summaries are directly comparable and so that trends across cycles can be identified by the reader.

The period covered section states the dates of the scan, the source types consulted, the volume of material reviewed, and any changes to the scan's scope or sources since the previous cycle. The signals identified section lists the signals that met the scan's inclusion criteria, grouped by driver or domain, with each signal's source, date, and a one-sentence statement of what the signal indicates. The emerging issues section presents the issues derived from clusters of signals, with each issue defined, supported by the signals that compose the cluster, and assessed for trajectory, uncertainty, and proximity. The assessment section ranks the emerging issues by their significance for the planning authority's questions and identifies those that warrant escalation to scenario construction or to the next foresight report.

The summary's standard is that every signal cited must be traceable to its source, every issue must be supported by a cluster of at least three signals, and every assessment must distinguish the evidence from the analyst's judgement. The summary is reviewed by a panel that includes the analysts who conducted the scan and at least one decision-maker who will use it. The published summary carries a date, a version number, and a statement of the next scheduled cycle. The summary feeds the foresight report when the assessment identifies issues that warrant escalation, and the report's findings section draws on the cumulative record of summaries produced since the previous report.

D.4 The Emerging Issues Brief Template

The emerging issues brief is the short document that presents a single emerging issue to decision-makers for a response. The template comprises five standard sections: the issue, the signals, the assessment, the implications, and the recommended response. The brief is typically four to eight pages, written for a non-specialist audience, and produced when an issue identified in the horizon scan summary has been judged to warrant a specific response. The format is fixed so that successive briefs on related issues are directly comparable and so that decision-makers can locate the information they need without re-learning the structure.

The issue section defines the emerging issue in a single paragraph, states the domain to which it belongs, and notes the period over which it has been observed. The signals section lists the signals that compose the issue's cluster, with each signal's source, date, and a one-sentence statement of what the signal indicates, and a note on any signals that point in a different direction. The assessment section presents the analyst's judgement on the issue's trajectory, uncertainty, and proximity, distinguishing the evidence from the judgement and noting the principal alternative interpretations. The implications section translates the issue into consequences for the plan's assumptions, risks, and revision triggers, identifying which decisions are affected and on what timescale.

The recommended response section sets out the action the planning authority should take, distinguishing between actions to be taken now and actions to be taken contingent on the issue's continued development. The brief's standard is that the recommendation must be supported by the assessment, the assessment must be supported by the signals, and the signals must be traceable to their sources. The brief is reviewed by a panel that includes the analyst who produced it and at least one decision-maker who will act on it. The published brief carries a date, a version number, and a statement of the next scheduled review, and the brief is filed alongside the horizon scan summary from which the issue was drawn.

D.5 Chapter Summary

The four templates set out in this appendix govern the principal foresight products. The foresight report is the comprehensive document that records an exercise's design, conduct, findings, and recommendations, with six standard sections populated to defined standards. The scenario briefing is the short document that presents a scenario set to decision-makers, with four standard sections and the requirement that each scenario be plausible, consistent, and distinct. The horizon scan summary is the periodic document that records the signals identified and the emerging issues derived, with four standard sections and the requirements of source traceability, cluster support, and evidence-judgement separation.

The fourth template, the emerging issues brief, presents a single issue for response, with five standard sections and the requirement that each layer of the brief support the one above it. Together the four templates constitute the standard product set of the foresight function. The planning authority that adopts them produces work that is comparable across cycles and auditable by reviewers. The templates share three properties that account for their authority: fixed structure, source traceability, and scheduled review.

Fixed structure makes successive products directly comparable and allows the reader to locate information without re-learning the format. Source traceability requires every signal, assessment, and recommendation to be traceable to its evidence. Scheduled review subjects every product to a panel that includes both producers and users and files the review record with the published product. The templates are published, dated, and version-controlled, and revisions are recorded so that products produced under earlier versions remain interpretable.

Appendix E — The Foresight Analyst's Reading List

This appendix provides an annotated reading list covering scenario methods, horizon scanning, expert elicitation, and comparative foresight practice. The list includes the foundational texts and the contemporary practice references.

E.1 Foundational Texts on Scenario Methods

The foundational texts on scenario methods establish the discipline's analytical lineage and the conceptual vocabulary subsequent practice has refined. Pierre Wack's two Harvard Business Review articles, "Scenarios: Uncharted Waters Ahead" and "Scenarios: Shooting the Rapids," published in 1985, set out the case that scenarios are instruments for shifting the mental models of decision-makers rather than forecasts of the future. Wack drew on his experience at Royal Dutch Shell to argue that the value of a scenario exercise lies in the reception of the scenarios by the strategy function, and that the exercise succeeds when decision-makers recognise futures they had previously dismissed. The contribution is the reframing of scenarios from an analytical product to a perceptual intervention, and this reframing is the foundation on which the subsequent literature rests.

Peter Schwartz's "The Art of the Long View," published in 1991, translated the practice developed in corporate settings into a methodological text accessible to a wider audience. Schwartz set out the structure of a scenario exercise that has become standard: identifying the focal question, enumerating the driving forces, assessing uncertainty and impact, selecting the scenario logics, and constructing the narratives. The text is the source of much of the discipline's vocabulary, including the terms "predetermined elements" and "key uncertainties," and its treatment of the role of intuition in scenario construction is the one most subsequent practitioners have built on. The contribution is the codification of the method into a sequence that can be taught, applied, and audited.

Kees van der Heijden's "Scenarios: The Art of Strategic Conversation," published in 1996, extended the Wack and Schwartz tradition by treating scenario construction as an instrument of organisational learning rather than a one-off analytical exercise. Van der Heijden, drawing on Royal Dutch Shell experience, argued that the scenario set is maintained over successive cycles and used as a standing frame against which the organisation tests its strategy. The text introduced the concept of the "business idea" and treated the scenario set as part of the institutional infrastructure of strategy. The contribution is the integration of scenarios into the ongoing strategic conversation of the organisation, and it is the text most directly responsible for the institutional framing the doctrine adopts.

E.2 Horizon Scanning and Weak Signals

Horizon scanning is the systematic detection of early indicators of change, and the foundational text on its analytical logic is Igor Ansoff's work on weak signals. Ansoff's 1975

article, “Managing Strategic Surprise by Response to Weak Signals,” and his 1984 book, “Implanting Strategic Management,” set out the proposition that strategic surprises typically announce themselves through signals too ambiguous to interpret at the time of their appearance. Ansoff proposed that organisations construct the capacity to detect such signals, assess their significance, and develop preparatory responses before the signals become unambiguous threats. The contribution is the analytical distinction between the signal and the threat, and the proposition that the gap between the two is the time available for preparation.

The United Kingdom’s Government Office for Science has published guidance on horizon scanning that translates Ansoff’s analytical logic into operational practice for government. The guidance sets out the components of a standing scan: the sources to be monitored, the scan protocol, the assessment of detected signals, the triage that separates weak signals from noise, and the products that convey the scan’s findings to the strategy function. The guidance is the most fully developed operational treatment of horizon scanning in a government context, and it has informed the design of scanning systems across several jurisdictions. The methodological contribution is the codification of the standing scan as a managed function, with the institutional infrastructure the function requires.

The European Commission’s Joint Research Centre, through its foresight publications, has contributed to horizon scanning methodology by integrating it with the broader foresight toolkit. The JRC’s work has treated horizon scanning as the input layer of the foresight cycle, feeding driver identification, scenario construction, and cross-impact analysis, and its publications have set out the protocols for the integration. The JRC has also contributed to scanning methodology through the European Forum on Forward Looking Activities and through its handbook on foresight methodology. The contribution is the integration of horizon scanning with the rest of the foresight cycle, and the treatment of scanning as one component of a coordinated analytical programme rather than a stand-alone activity.

E.3 Expert Elicitation

Expert elicitation is the systematic collection and aggregation of expert judgment, and the foundational text on its most widely used method is the RAND Corporation’s Delphi publications. The Delphi method, developed at RAND in the 1950s by Norman Dalkey and Olaf Helmer, used iterative anonymous questionnaires and controlled feedback to aggregate the judgments of a panel of experts. The RAND publications set out the rationale for anonymity, the structure of the iterative rounds, and the aggregation rules for the final estimates, and the method has been adapted to foresight, technology assessment, and policy analysis. The contribution is the formalisation of expert judgment into a method that can be replicated, documented, and audited, and that mitigates the biases that arise in face-to-face expert gatherings.

The literature on structured expert judgment extends the Delphi tradition by treating the elicitation procedure as the object of methodological design. Roger Cooke’s “Experts in Uncertainty,” published in 1991, set out the classical model for structured expert judgment, in which experts’ assessments are weighted by their performance on calibration questions drawn from the relevant domain. The classical model is the most widely used approach to structured expert

judgment in policy contexts, and its applications have been documented in risk assessment, environmental modelling, and engineering reliability. The methodological contribution is the introduction of performance-based weighting, which addresses the problem that experts vary in their accuracy and that equal weighting of expert judgments can produce aggregated estimates worse than those of the best individual expert.

The literature on bias and debiasing in expert judgment provides the empirical foundation for the procedural design of elicitation. Daniel Kahneman's work, summarised in "Thinking, Fast and Slow" (2011), and the earlier research conducted with Amos Tversky, set out the systematic biases to which expert judgment is subject, including overconfidence, anchoring, availability, and the substitution of an easier question for the one posed. The literature on debiasing, developed across psychology, behavioural economics, and decision analysis, sets out the procedural interventions that mitigate these biases: structured question framing, the elicitation of ranges rather than point estimates, the use of multiple experts, and the documentation of reasoning. The methodological contribution is the empirical grounding of elicitation procedure in the documented characteristics of expert judgment, and the proposition that the procedure is the principal instrument for managing the quality of the elicited judgment.

E.4 Comparative Foresight Practice

Comparative foresight practice is the literature on how foresight is conducted across mature jurisdictions, and the principal source is the OECD's work on strategic foresight. The OECD's International Futures Programme, conducted from the 1990s, and the more recent work under the OECD Future Programme, have produced comparative analyses of foresight institutions, methods, and applications across member states. The OECD publications are the most comprehensive comparative source available, and they cover both the institutional arrangements and the methodological practice of the jurisdictions surveyed. The contribution is the comparative frame that allows a jurisdiction to assess its own practice against that of its peers, and the documentation of the range of institutional forms the practice takes.

The publications of Singapore's Centre for Strategic Futures provide the most developed documentation of foresight practice in a single jurisdiction. The CSF's publications, including its methodological notes on scenario construction and on the Cyclic Early Warning system, set out how foresight is conducted within the Singapore administrative tradition. The CSF's contribution to the comparative literature is the documentation of the placement of the foresight function within the central coordination machinery of government, and the demonstration of how that placement shapes the function's products and its relationship to the strategy function. The publications are also a source on the application of foresight to whole-of-government planning, an application less developed in the literature of other jurisdictions.

The reports of Finland's Committee for the Future and the publications of France Stratégie complete the comparative picture for European jurisdictions. Finland's Committee for the Future, established in 1993, produces reports on long-term issues for the Parliament, and its publications document the legislative engagement with foresight that distinguishes the Finnish model. France Stratégie, the French government's strategic planning commission, publishes analyses on long-term

issues across economic, social, environmental, and technological policy, and its publications document the independent advisory model in which the commission's authority derives from its analytical standing rather than from line responsibility. Together with the OECD and the Singaporean sources, these publications constitute the comparative literature a foresight practitioner should read to understand the range of institutional forms and the principles that underlie them.

E.5 The Planning Manual

The Planning Manual is the general doctrine of which the doctrine is the specialty extension, and the relationship between the two is the relationship of general to specialty doctrine. The Planning Manual sets out the general doctrine of planning: the planning cycle, the institutional arrangements, the analytical methods planning draws on, and the integration of planning with the budget and with political authority. The Handbook extends that doctrine on the foresight specialty, treating horizon scanning, scenario construction, expert elicitation, and the institutionalisation of foresight at the depth the specialty requires. The planner should read the two together, because the Handbook's specialty doctrine is incomplete without the Manual's general doctrine, and the Manual's general doctrine does not address the foresight specialty at the depth required.

The points of connection between the Handbook and the Manual are the points at which the specialty doctrine extends the general. The Planning Manual's treatment of the planning cycle is the framework within which the Handbook's treatment of the foresight cycle is set, and the Manual's treatment of institutional arrangements is the framework within which the Handbook's treatment of the foresight unit is set. Where the Planning Manual addresses a topic adequately, the handbook references it and goes deeper on the specialty-specific craft. Where the Manual does not address a topic because the topic is specialty-specific, the Handbook sets out the doctrine on its own.

The order of reading follows the relationship. A planner new to the Foresight Handbook should read the Planning Manual's chapters on the planning cycle, on institutional arrangements, and on analytical methods, then read the Handbook's chapters on the definition of foresight, on the foresight cycle, on the methods, and on institutionalisation. The reading order matters because the Handbook's specialty doctrine assumes the Manual's general doctrine, and the planner who reads the Handbook without the Manual will find the doctrine undercontextualised. A planner revising the Handbook's doctrine should likewise revisit the Manual, because revisions to the general doctrine have consequences for the specialty doctrine, and revisions to the specialty doctrine must remain consistent with the general.

E.6 Chapter Summary

Appendix E has set out an annotated reading list covering the foundational texts and the contemporary practice references the foresight practitioner should consult. The foundational texts on scenario methods are Wack's two Harvard Business Review articles, Schwartz's "The Art of the Long View," and van der Heijden's "Scenarios: The Art of Strategic Conversation," which together

established the discipline’s analytical lineage and conceptual vocabulary. The readings on horizon scanning and weak signals include Ansoff on weak signals, the United Kingdom’s Government Office for Science guidance, and the European Commission’s Joint Research Centre publications, which together develop the analytical logic and the operational practice of scanning.

The readings on expert elicitation include the RAND Corporation’s Delphi publications, the literature on structured expert judgment represented by Cooke’s classical model, and the work on bias and debiasing in expert judgment represented by Kahneman and Tversky’s research. The readings on comparative practice include the OECD’s work on strategic foresight, the publications of Singapore’s Centre for Strategic Futures, the reports of Finland’s Committee for the Future, and the publications of France Stratégie, which together constitute the comparative literature on institutional forms and practice. The Planning Manual is the general doctrine of which the doctrine is the specialty extension, and the planner should read the two together, in the order general doctrine then specialty doctrine.