

The Tychevia Process

Julian Macnamara
(Glandore Associates)

April 2026

Abstract

Most approaches to method in policy, organisational change, and strategy take the form of toolkits: collections of techniques to be selected and applied according to context. While such approaches can be effective in conditions of stability, they frequently fail in environments characterised by uncertainty, contestation, and systemic instability.

This failure is not primarily technical, but epistemic. Techniques are applied before the nature of the problem is properly understood, evidence is accumulated without establishing relevance, and solutions are proposed before the system in which they must operate is made visible. The result is premature convergence, false certainty, and recurrent reform that addresses symptoms while leaving underlying dynamics untouched.

This paper introduces the Tychevia Process as a structured alternative. Rather than presenting methods as interchangeable tools, Tychevia treats inquiry as a disciplined sequence of thinking, in which different techniques become necessary at different stages as understanding evolves. Beginning with an epistemically grounded question, the process proceeds through orientation, sense-preparatory research, the construction of a Knowledge Engine, and iterative sense-making using lenses, causal modelling, and root convergence, before arriving at proposals that are explicitly conditioned by uncertainty and constraint.

The contribution of the paper is twofold. First, it makes explicit the ordering and rationale that underpins effective use of analytical techniques in complex contexts. Second, it offers a coherent framework within which additional methods can be incorporated over time without encouraging premature solutionism or loss of epistemic discipline.

In this account, method is not a means of control or optimisation, but a way of maintaining coherent judgement in situations where certainty is unavailable and action must proceed without illusion.

1	Introduction	3
2	The Epistemic Starting Point: The Question	3
3	Orientation and Problem Typing	4
4	Desk Research as Sense-Preparation	5
5	The Knowledge Engine	6
5.1	The Static Knowledge Engine	6
5.2	The Active Knowledge Engine	6
6	Sense-Making as an Iterative Discipline	7
6.1	Wicked structures and systemic instability	8
6.2	Lenses and relevance	8
6.3	A concrete example: the governance and power lens	8
6.4	Causal structure and dynamics	10
6.5	Root convergence and drilling methods	10
6.6	“5 Whys” in context	11
6.7	Contrast with Fishbone (Ishikawa) diagrams	12
6.8	Dialectical synthesis and the work of Georg Hegel	13
7	From Sense-Making to Proposals	14
8	Approaches to implementing proposals	15
8.1	Work-Out-style interventions	15
8.2	Scaled innovation forums	15
8.3	Scenario-based evaluation	16

1. Introduction

Most approaches to method in policy, organisational change, and strategy take the form of toolkits: collections of techniques to be selected, combined, or applied according to taste. In conditions of stability, such catalogues can be useful. In conditions of uncertainty, contestation, and systemic instability, they frequently fail.

The failure is not primarily technical. It is epistemic. Techniques are deployed before the nature of the problem is properly understood, evidence is accumulated before relevance is established, and solutions are proposed before the system in which they must operate is made visible. The result is premature convergence, false certainty, and recurrent reform that addresses symptoms while leaving underlying dynamics untouched.

Tychevia was developed in response to this failure mode. It does not present methods as interchangeable tools, nor does it privilege any single analytical frame. Instead, it treats method as a disciplined sequence of thinking, in which different techniques become necessary at different moments as understanding evolves. The order in which techniques are applied is therefore not incidental; it is constitutive of the approach.

This paper describes the technique path that underpins Tychevia. It sets out how inquiry proceeds from an initial question, through orientation, desk research, and the creation of a knowledge engine, into iterative sense-making and, finally, the formulation of proposals. The emphasis throughout is not on the mechanics of individual techniques, but on why they are used when they are, and what goes wrong when the sequence is disrupted.

The intention is twofold. First, to make explicit the epistemic discipline that differentiates Tychevia from method catalogues and best-practice frameworks. Second, to provide a structure that allows additional techniques to be incorporated over time without undermining coherence or encouraging premature solutionism.

In what follows, method is treated not as a means of control or optimisation, but as a way of maintaining coherent judgement in situations where certainty is unavailable and good thinking is often crowded out by speed, narrative pressure, and institutional habit.

2. The Epistemic Starting Point: The Question

Every Tychevia inquiry begins with a question, but not in the sense commonly implied by problem statements, briefs, or terms of reference. The question is not a request for an answer. It is an epistemic commitment: a declaration of what must be understood well enough to act without false certainty.

In complex and contested contexts, poorly formed questions are one of the primary sources of failure. They smuggle in assumptions about causality, ownership, solvability, or success before inquiry has even begun. Questions framed in terms of delivery, optimisation, or solution selection

tend to collapse uncertainty prematurely and channel subsequent analysis toward confirmation rather than understanding.

The role of the initial question in Tychevia is therefore orienting rather than directive. It establishes the scope of inquiry, signals what is not yet known, and resists the temptation to define success in advance. A well-formed question makes explicit what kind of knowing is required and what kinds of certainty are unavailable or inappropriate.

Examples of such questions include:

- ❖ What is this system repeatedly failing to resolve, despite sustained effort?
- ❖ Where does responsibility exist without corresponding authority?
- ❖ Why do similar interventions produce different outcomes across time or context?

These questions do not point directly to action. They create the conditions under which orientation, sense-making, and disciplined analysis can occur without being captured by existing narratives or institutional habit.

Importantly, the question is not fixed once posed. As inquiry progresses, questions may be refined, narrowed, or re-articulated. However, this evolution is deliberate rather than reactive. Changes to the question reflect increased understanding of the system, not impatience with uncertainty.

By treating the question as the epistemic starting point rather than a prelude to technique, Tychevia establishes a foundation on which subsequent methods can be applied coherently and in the appropriate sequence.

3. Orientation and Problem Typing

Before analysis begins, Tychevia places deliberate emphasis on orientation. This step asks a deceptively simple question: what kind of problem are we actually dealing with? The answer determines not only which techniques are appropriate, but which forms of certainty are available and which are not.

A common failure mode in policy and organisational work is category error: treating complex or wicked situations as if they were merely complicated, or assuming that increased data or optimisation will resolve issues that are structurally contested. Orientation exists to prevent this error from becoming embedded in subsequent analysis.

Tychevia distinguishes between problems that are broadly tame, complicated, complex, and wicked. These categories are not labels of difficulty, but of structure. Tame problems admit clear definitions and repeatable solutions. Complicated problems may be technically demanding but are ultimately decomposable. Complex problems involve adaptation, feedback, and non-linearity. **Wicked problems** are characterised by unstable definitions, contested ownership, moral load, and the absence of stopping rules.

In practice, many real-world situations combine elements of more than one category. Orientation work therefore focuses less on classification than on dominance: which characteristics shape the system's behaviour most strongly, and where attempts at optimisation or control are likely to misfire.

Orientation also involves boundary setting. Decisions are made about what is in scope, what is out of scope, and what cannot be influenced directly. These boundaries are treated as provisional but explicit. Making them visible reduces the risk of implicit assumptions driving later conclusions.

By establishing problem type and boundaries early, Tychevia creates a shared understanding of what inquiry can reasonably achieve. This does not simplify the problem, but it does prevent techniques from being applied in ways that generate confidence without insight.

Only once orientation is complete does it become meaningful to engage in desk research or structured sense-making. Without this step, even rigorous analysis risks reinforcing the very misunderstandings it seeks to resolve.

4. Desk Research as Sense-Preparation

Desk research in Tychevia is deliberately restrained. Its purpose is not to establish authority through accumulation, nor to construct a definitive account of the system. Instead, it functions as sense-preparation: creating the conditions under which meaningful inquiry can occur without being captured by existing narratives.

In complex and wicked contexts, desk research is often overused and misapplied. Reports, data sets, evaluations, and reviews proliferate, yet understanding does not deepen proportionally. The risk is not ignorance but saturation — an excess of material that obscures rather than clarifies what is contested, assumed, or structurally unresolved.

Tychevia approaches desk research with a different intent. Existing material is examined to surface assumptions, identify points of disagreement, and recover fragments of institutional memory. Particular attention is paid to how problems have been framed over time, which explanations recur, and which lines of inquiry have quietly disappeared. The aim is not to decide which account is correct, but to understand why certain accounts persist and others do not.

Proportionality is central. The depth and breadth of desk research are shaped by the orientation work that precedes it. Where uncertainty is structural, further evidence does not resolve it; where history is relevant, forgetting it is actively harmful. Desk research therefore stops when it has achieved its purpose: clarifying what is known, what is contested, and what remains opaque.

Crucially, desk research does not yet constitute analysis. Its outputs are not conclusions but inputs. They inform the creation of a Knowledge Engine by defining the material, perspectives, and histories that must be held in view as sense-making begins.

By treating desk research as preparation rather than proof, Tychevia avoids a common trap: mistaking familiarity with the literature for understanding of the system. Only once this groundwork is laid does it become possible to move from accumulated knowledge to structured inquiry.

5. The Knowledge Engine

The transition from desk research to sense-making is not automatic. Accumulated material, however carefully curated, does not yet constitute understanding. Tychevia addresses this gap through the creation of a Knowledge Engine: a structured epistemic container within which inquiry can proceed coherently.

The Knowledge Engine is not a database, a model, or a repository. It is a deliberate arrangement of sources, perspectives, and modes of interrogation that defines what knowledge is brought into play and how it is worked with. Its purpose is to prevent analysis from drifting, collapsing into a single frame, or being dominated by the loudest narrative.

In practice, the Knowledge Engine has two interrelated components: a static element and an active element. Together, they establish both the content of inquiry and the discipline by which that content is examined.

5.1. The Static Knowledge Engine

The static component defines what material is held in view. This typically includes selected documents, data sources, historical records, prior analyses, and formal accounts relevant to the Domain of Interest. Inclusion is purposeful rather than exhaustive. Material is chosen because it shapes how the problem has been understood, governed, or acted upon, not because it is comprehensive or authoritative.

Importantly, the static Knowledge Engine preserves disagreement. Competing explanations, unresolved evaluations, and contradictory evidence are retained rather than reconciled. This ensures that subsequent analysis does not begin from an artificial consensus.

5.2. The Active Knowledge Engine

The active component defines how inquiry is conducted. It brings together analytical lenses, structured dialogue, iterative questioning, and, where appropriate, multiple agents or perspectives. This is the mechanism through which the static material is interrogated, reinterpreted, and tested.

The active Knowledge Engine resists linear analysis. Questions are revisited, assumptions are challenged, and provisional interpretations are allowed to evolve. What matters is not speed of convergence, but the quality of insight generated and the avoidance of premature closure.

By separating what knowledge is included from how it is examined, the Knowledge Engine creates a stable foundation for sense-making without fixing conclusions in advance. Techniques such as causal modelling, root analysis, or dialectical reasoning can now be applied in sequence, each drawing on a shared epistemic base.

This structural move is subtle but consequential. Without a Knowledge Engine, methods float free and compete for primacy. With it, they become coordinated instruments within a single inquiry, summoned when the evolving understanding of the system makes them necessary.

6. Sense-Making as an Iterative Discipline

Sense-making in Tychevia is not a single analytical step, nor a sequence of techniques applied once and set aside. It is an iterative discipline through which understanding is progressively refined while remaining open to revision. The aim is not to produce a definitive explanation of the system, but to make its structure, tensions, and recurring patterns visible enough to support coherent judgement.

This phase begins only once orientation is complete and the Knowledge Engine has been established. Without those foundations, sense-making risks oscillating between abstraction and detail without ever stabilising insight.

6.1. Wicked structures and systemic instability

Sense-making typically starts with the identification of wicked structures rather than with problem statements. Attention is paid to instability of definition, contested ownership, moral load, and recurrence across reform cycles. These characteristics signal that the system cannot be optimised in any simple sense and that conventional notions of solution are likely to mislead.

The purpose at this stage is not to label a problem as wicked, but to recognise where instability is structural and where attempts at resolution are likely to generate displacement rather than improvement.

6.2. Lenses and relevance

Analytical lenses are then applied to examine the same problem space from deliberately contrasting standpoints. Lenses do not provide answers; they redefine what counts as relevant. By holding the problem constant and changing the interpretive frame, blind spots, tensions, and competing explanations are surfaced.

Tychevia employs a small, disciplined set of lenses, typically including systems, governance and power, incentives, culture and narrative, lived experience, and time or institutional memory. Agreement between lenses is treated with caution. Productive disagreement is often more informative than rapid convergence.

6.3. A concrete example: the governance and power lens

To illustrate how lenses are used in practice, consider the governance and power lens. This lens examines how authority, accountability, and decision rights are distributed within a system, and where formal responsibility diverges from actual capacity to act.

The lens asks questions such as:

- ❖ Who has the authority to make decisions that materially affect outcomes?
- ❖ Who is held accountable when those outcomes fall short?
- ❖ Where are decisions escalated, deferred, or displaced?
- ❖ Which actors carry risk without corresponding control?

When applied to a public sector health system, the governance and power lens immediately shifts attention away from performance metrics or operational efficiency and toward structural tensions. For example, repeated service failures may initially be framed as issues of capacity, capability, or local management. Through the governance lens, a different pattern often emerges: responsibility for outcomes is devolved to local organisations, while strategic control, funding decisions, and political accountability remain centralised.

This reframing does not yet explain why failures recur, nor does it suggest a solution. Its purpose is narrower and more disciplined. It makes visible a structural condition — responsibility without authority — that would remain obscured if the problem were viewed solely through operational or financial frames.

Once this condition has been surfaced, other methods can be applied meaningfully. **Causal Loop Diagrams** can explore how governance asymmetries interact with demand, workforce behaviour, and political cycles. Drilling methods such as “5 Whys” can then be used to examine why attempts to correct performance repeatedly stall or reverse. In some cases, dialectical reasoning becomes necessary to work with persistent tensions between local autonomy and central control rather than attempting to resolve them prematurely.

Crucially, the governance and power lens does not claim primacy. When applied alongside other lenses — such as incentives, culture, or lived experience — it may be challenged or complemented. Its value lies in what it makes visible, not in asserting that governance is the sole or dominant explanation.

This example illustrates the role of lenses in Tychevia more generally. Lenses do not diagnose problems or prescribe action. They alter relevance. By changing what is foregrounded and backgrounded, they shape the subsequent use of causal modelling, root convergence, and synthesis without collapsing inquiry into a single explanatory frame.

6.4. Causal structure and dynamics

Insights generated through lenses are then integrated using causal modelling, most commonly through Causal Loop Diagrams. These diagrams are not predictive models. They are sense-making artefacts that reveal feedback loops, delays, reinforcing and balancing dynamics, and points of fragility or lock-in.

Placing causal structure at this stage is critical. It allows patterns to be understood dynamically rather than descriptively and provides a shared reference point for subsequent inquiry.

6.5. Root convergence and drilling methods

Root convergence refers to the disciplined exploration of why particular patterns, failures, or tensions reappear over time, despite repeated intervention. Its purpose is not to identify a single, definitive cause, but to distinguish between surface manifestations and systemic drivers that persist across contexts.

In Tychevia, root convergence is always preceded by orientation, lens-based sense-making, and causal structuring. This sequencing is critical. Drilling methods deepen inquiry within an established frame; they do not define relevance themselves. When applied too early, they produce convincing but misleading explanations.

A range of techniques may be appropriate at this stage, depending on the nature of the system and the question being explored. These include structured post-mortems, cross-case comparison, temporal analysis of prior reforms, and constraint mapping. Each of these approaches can reveal recurring conditions or decision patterns that are not visible at the level of individual events.

However, the most commonly used drilling technique within Tychevia is “5 Whys”, applied with care and often in parallel.

6.6. “5 Whys” in context

“5 Whys” is a simple but powerful method for exploring causality by repeatedly asking why a particular outcome or pattern occurs. Each iteration seeks to move beyond immediate explanations toward deeper conditions that make the outcome likely or persistent.

In its simplest form, the method proceeds as follows:

- ❖ A pattern or issue is identified.
- ❖ The question “Why does this occur?” is asked.
- ❖ The answer is treated as provisional, and the question is asked again.
- ❖ This process is repeated until further questioning no longer produces materially different insight.

Despite its apparent simplicity, “5 Whys” is frequently misunderstood. It is not a technique for uncovering a single root cause, nor does it guarantee correctness through repetition. Its value lies in exposing assumptions about causality and in revealing how explanations change as inquiry deepens.

In Tychevia, “5 Whys” is never used in isolation. It is applied within an explicit lens and often conducted multiple times from different standpoints. For example, a governance lens may produce a different chain of reasoning from an incentives or cultural lens. Convergence across these parallel drills is treated as significant; divergence is treated as informative rather than problematic.

This multi-perspective application guards against one of the common failure modes of “5 Whys”: premature convergence on an explanation that reflects the initial framing more than the system itself. By comparing drilling paths, Tychevia surfaces echo patterns — conditions or dynamics that recur regardless of the lens applied.

Importantly, the outcome of “5 Whys” is not a recommendation. It is an enriched understanding of why certain dynamics persist. This understanding informs subsequent judgement about where intervention is possible and where restraint is warranted.

Root convergence concludes not when a cause has been identified, but when the system's propensity to reproduce particular outcomes has been sufficiently illuminated to allow action without illusion.

6.7. Contrast with Fishbone (Ishikawa) diagrams

Fishbone, or Ishikawa, diagrams are a well-established technique for organising potential causes of a defined problem into categories such as process, people, technology, environment, or policy. They are particularly effective in situations where the problem statement is stable and the aim is to ensure that contributory factors are not overlooked.

In such contexts, Fishbone diagrams serve as a useful completeness check. They encourage systematic consideration of multiple cause categories and can support team-based diagnosis where the primary risk is omission rather than misframing.

However, Fishbone diagrams embed assumptions that limit their usefulness in complex and wicked systems. Most notably, they assume that the problem itself is sufficiently well-defined to act as a fixed reference point, and that causes can be meaningfully enumerated and arranged around it. This makes them well suited to failure analysis in bounded processes, but less effective where problem definitions are contested, dynamic, or politically loaded.

In Tychevia's work, these assumptions often do not hold. Problems are frequently unstable, and causal explanations evolve as inquiry proceeds. In such cases, the visual completeness of a Fishbone diagram can be misleading, creating the appearance of thoroughness while obscuring deeper structural dynamics and feedback effects.

For this reason, Tychevia tends to prefer drilling methods such as 5 Whys, applied after relevance has been established through lenses and causal modelling. Where Fishbone diagrams ask "what might be contributing to this problem?", drilling methods ask "why does this pattern keep reappearing, even as interventions change?". The latter question aligns more closely with the aim of root convergence in systems characterised by recurrence and adaptation.

This is not to suggest that Fishbone diagrams are inappropriate in all circumstances. Rather, their use is contingent on problem type. Where problems are bounded, definitions are stable, and learning objectives are clear, they may play a supporting role. Where these conditions are absent, Tychevia deliberately avoids techniques that risk premature closure or false reassurance.

6.8. Dialectical synthesis and the work of Georg Hegel

The term dialectical synthesis originates in the work of the nineteenth-century philosopher Georg Wilhelm Friedrich Hegel. Hegel was concerned not with formal logic or argumentation, but with how understanding develops over time when reality presents persistent contradictions that cannot be resolved by choosing one side over another. His central insight was that certain conflicts are not errors to be eliminated, but signals that existing ways of thinking are insufficient.

In its simplest form, Hegelian dialectic is often summarised as thesis, antithesis, and synthesis. While this shorthand is imperfect, it captures the essential movement. A thesis represents a dominant way of framing a situation. An antithesis represents a competing or opposing framing that cannot be dismissed. A synthesis does not compromise between the two, nor does it select a winner. Instead, it reframes the problem so that the original opposition is no longer the primary organising tension.

In Tychevia, dialectical synthesis is not used as a philosophical lens or as a method of analysis in its own right. It is employed selectively, and only when sense-making reveals tensions that repeatedly block progress. It becomes relevant when multiple explanations are simultaneously compelling, interventions oscillate between alternatives, or debates recycle without resolution.

A concrete example helps clarify this use.

Consider a health system repeatedly caught between the imperatives of central control and local autonomy. Centralisation is defended as necessary for equity, consistency, and fiscal discipline. Local autonomy is defended as essential for responsiveness, professional judgement, and adaptation to local conditions. Attempts to resolve the tension by swinging from one position to the other often result in reform cycles rather than improvement.

Dialectical synthesis does not ask which position is correct. Instead, it asks what is trying to emerge through the conflict. In this case, the synthesis may involve reframing the problem around stewardship rather than control or autonomy: clarifying which decisions must be held centrally to protect system integrity, and which must be decentralised to enable learning and adaptation. The original opposition is not eliminated, but it is reorganised within a more generative frame.

Within Tychevia's technique path, dialectical synthesis typically follows root convergence. Lenses and drilling methods surface the persistence of the tension and the reasons it recurs. Dialectical reasoning then provides a way of moving forward without pretending that the tension can be resolved through optimisation or technical fixes.

Importantly, dialectical synthesis does not guarantee agreement, nor does it produce solutions in the conventional sense. Its contribution lies in enabling reframing where binary choices have

repeatedly failed, allowing subsequent proposals to be shaped by a more adequate understanding of the system's underlying contradictions.

7. From Sense-Making to Proposals

In Tychevia, proposals are not the objective of inquiry, but its consequence. They emerge only once the structure of the problem has been made sufficiently visible to allow action without illusion. This distinguishes proposals from solutions, recommendations, or implementation plans, all of which assume a degree of certainty that is rarely available in complex or wicked systems.

The transition from sense-making to proposals is therefore deliberate. Lenses, causal modelling, root convergence, and, where necessary, dialectical synthesis are used to reduce the risk of acting on misframed problems. Only when these techniques have clarified what is driving recurrence, where constraints genuinely lie, and which tensions cannot be resolved does it become responsible to consider intervention.

Proposals in Tychevia typically take the form of structured options rather than prescriptions. They articulate possible pathways for action, the conditions under which those pathways might be viable, and the trade-offs they entail. Importantly, they also make explicit what cannot be achieved, at least in the short term, and where restraint is preferable to activity.

In public sector and NHS contexts, proposals often function as holding structures rather than programmes of change. They are designed to stabilise decision-making, protect institutional memory, and prevent further degradation of trust or capability. Success is judged less by immediate outcomes than by whether the system's capacity for coherent judgement and learning is preserved.

In corporate and financial contexts, proposals may be more interventionist, but the underlying discipline remains the same. Actions are shaped by an explicit understanding of uncertainty, narrative pressure, and systemic feedback, rather than by optimisation against a single performance metric.

Across all domains, Tychevia proposals share three characteristics. First, they are grounded in an explicit account of why the problem persists. Second, they respect structural constraints rather than attempting to wish them away. Third, they are designed to remain intelligible over time, even as conditions change.

This orientation reflects a broader commitment within Tychevia: to treat action as provisional, learning as cumulative, and judgement as something to be exercised and protected rather than automated. In this sense, proposals mark not the end of inquiry, but the point at which understanding is sufficient to act responsibly.

8. Approaches to implementing proposals

Tychevia is concerned with problem definition, sense-making, and the articulation of proposals. It does not prescribe how proposals should be implemented, nor does it assume authority over their selection. Responsibility for deciding which proposals to pursue, and how to pursue them, rests explicitly with the project stakeholder or originator, who remains accountable for implementation choices and outcomes.

Once proposals have been developed and understood in the context of a stabilised causal model, a range of implementation approaches may be employed. The choice of approach is contextual and should reflect the nature of the proposal, the level of risk and reversibility involved, and the degree of coordination, consent, or experimentation required.

8.1. Work-Out-style interventions

Work-Outs are structured, time-bounded interventions designed to convert agreed proposals into concrete commitments and actions. Originating at General Electric, Work-Outs bring together cross-boundary actors to resolve interdependencies, remove organisational blockages, and assign ownership with clear timelines.

In this context, a Work-Out does not revisit problem definition or proposal selection. Instead, it functions as an action-forcing mechanism, used where authority is sufficiently clear but execution is impeded by fragmentation, delay, or ambiguity in decision rights.

8.2. Scaled innovation forums

Scaled innovation forums, such as **IBM's Innovation Jam**, are digitally mediated, time-limited events designed to engage large and diverse groups of participants in the exploration and refinement of a defined proposal or implementation challenge. These forums can be used to surface implementation variants, identify risks or unintended consequences, and build shared understanding or legitimacy across organisational boundaries.

Unlike open-ended ideation exercises, the use of innovation forums in this framework is deliberately constrained. The underlying proposal is treated as given; the forum explores how it might be implemented, adapted, or supported in different contexts. Final selection and investment decisions remain with accountable stakeholders.

8.3. Scenario-based evaluation

In systems characterised by tight coupling, high uncertainty, or significant second-order effects, proposals may be evaluated through structured scenario-based exploration prior to execution. In this mode, a wicked problem expressed as a causal loop diagram is provisionally frozen and used as a working model for “what if” analysis.

Proposed interventions are introduced hypothetically, and their potential consequences, feedback effects, and trade-offs are examined across alternative scenarios. Scenario-based evaluation does not seek prediction or optimisation; rather, it provides a disciplined way to rehearse action, refine proposals, and clarify risks before irreversible commitments are made.

These approaches are not sequential stages, nor are they mandatory. They represent alternative implementation pathways that may be selected, combined, or omitted according to context. In all cases, Tychevia provides the epistemic grounding against which implementation choices can be made deliberately and transparently.