

27. When Flexibility Replaces Explanation

Forecasting cosmology in the era of generational surveys

Ian Salisbury, Oxford, 22 December 2025

[Go to the catalogue](#)

Abstract

Forthcoming generational surveys such as the Nancy Grace Roman Space Telescope, alongside Euclid, Rubin, and DESI, are increasingly justified and optimised through forecast-driven analyses conducted well in advance of data acquisition. These forecasts play an essential operational role, but they also do something more consequential: they define, implicitly, what counts as cosmological success. Scientific value is routinely expressed in terms of constraining power over extended and deformable parameter spaces, with complementarity measured by the efficient reduction of degeneracies. This commentary argues that such practices, while methodologically effective and institutionally rational, risk shifting explanatory standards from structural necessity toward parametric accommodation. Using void-based forecasts as a representative case, it examines how observables that probe large-scale geometry directly are

nonetheless subsumed within highly flexible interpretive frameworks. The analysis places this tendency within the institutional coupling of investors, survey designers, and forecast metrics, showing how flexibility is structurally rewarded while hard constraints struggle to find methodological footing. In contrast, geometry-first approaches such as Penrose's Weyl curvature hypothesis and the isotropic singularity models of Ellis and Tod exemplify an alternative explanatory posture in which admissibility is enforced rather than negotiated. The Time-Symmetric Parity Constraint is briefly noted as a recent example of an attempt to pursue explanation through a hard geometric restriction. The aim is not to adjudicate between models, but to ask whether, at a moment of unprecedented observational precision, cosmology risks allowing flexibility to replace explanation.

1. Introduction: why this moment matters

Cosmology is entering a generational phase. Surveys such as Euclid,¹ Rubin,² DESI,³ and the Nancy Grace Roman Space Telescope⁴ are not incremental extensions of earlier

¹ *Euclid*: ESA medium-class mission with a reported cost of approximately €1.4–1.6 billion ($\approx$ US \$1.5–1.7 billion), excluding some nationally funded instrument contributions. Euclid is designed to map the large-scale distribution of galaxies and weak gravitational lensing over much of the extragalactic sky, with the primary aim of constraining dark energy, dark matter, and cosmic geometry.

² *Vera C. Rubin Observatory*: Construction and commissioning costs of approximately US \$0.7–0.8 billion, excluding full survey operations over the planned ten-year lifetime. Rubin will conduct the Legacy Survey of Space and Time, a wide, deep, time-domain imaging survey intended to probe cosmic structure, gravitational lensing, and transient phenomena.

³ *DESI*: Ground-based spectroscopic survey with an estimated construction cost of US \$100–150 million, funded primarily by the US Department of Energy with international partners. DESI is designed to obtain precise redshifts for tens of millions of galaxies and quasars in order to measure the expansion history of the universe and the growth of large-scale structure.

⁴ *Roman Space Telescope*: NASA flagship mission with a total lifecycle cost of approximately US \$4 billion, including spacecraft, instruments, launch, and initial operations. The figure is conservative and reflects post-rebaselining estimates. Roman is designed to conduct wide-field infrared imaging and spectroscopy to study large-scale structure, dark energy, and gravitational lensing with high statistical power.

programmes but long-horizon commitments that will shape observational practice and interpretation for decades. Increasingly, the scientific return of these surveys is articulated in advance through detailed forecasts that estimate their expected constraining power on cosmological parameters. These exercises play an essential operational role. They guide instrument design, survey strategy, and resource allocation. They also do something more consequential: they preconfigure the frameworks within which future data will be understood.

Forecasts are indispensable for planning surveys. Their growing role means that the language of “success” is often set in advance, in terms of parameter constraints and the reduction of degeneracies. This commentary asks whether that practical framing, repeated across major projects, is quietly changing what cosmology counts as explanation.

Historically, cosmological explanation relied on a balance between empirical adequacy and structural constraint. Geometric assumptions, symmetry principles, and boundary conditions sharply limited the space of admissible models, and explanatory success consisted in showing why the universe must behave in a particular way given those restrictions. Parameter freedom was introduced sparingly and provisionally, as a means of testing whether such constraints were too severe. In the contemporary forecasting regime, by contrast, flexibility is often treated as a baseline virtue. Models are valued for their capacity to accommodate a wide range of behaviours smoothly, and surveys are optimised accordingly.

This shift is not an error. It reflects the extraordinary precision now achievable in observational cosmology and the practical need to manage complex datasets. However, it raises a legitimate methodological question. When interpretive frameworks are defined primarily by their ability to absorb discrepancies through adjustable parameters, what becomes of explanation as a constraint-driven enterprise? At what point does preserving concordance through flexibility risk displacing the very notion of necessity that cosmological explanation traditionally sought?

These questions are not directed at any particular experiment or analysis. They arise from the cumulative logic of the field at a moment when observational power is unprecedented and interpretive commitments are increasingly locked in before data arrive. The aim of this commentary is to make that logic explicit and to examine its implications. By situating current forecast practices within a broader history of cosmological explanation, and by contrasting them with geometry-first approaches that deliberately restrict admissible models, the discussion seeks to clarify what is at stake. The issue is not which model is correct, but what standards of explanation cosmology now wishes to uphold.

2. Parameter freedom as a methodological norm

The expansion of cosmological parameter space has been gradual, pragmatic, and largely uncontroversial at each step. The introduction of the cosmological constant, followed by its reinterpretation as a dark-energy component, represented an early relaxation of otherwise rigid dynamical expectations. Subsequent extensions, including constant- w and time-varying w parameterisations, were initially framed as exploratory tools: provisional degrees of freedom introduced to test whether observational tensions genuinely required new physics. Over time, however, these extensions have migrated from the periphery to the centre of cosmological practice.

In contemporary forecasting studies, flexible parameterisations are no longer treated as hypotheses under test but as default elements of the model space. Survey performance is routinely assessed in terms of how efficiently such parameters can be constrained, and optimisation strategies are designed accordingly. This shift is subtle but significant. It marks a transition from using parameters to probe the limits of a framework to using them to preserve its viability across an expanding range of observations.

This methodological evolution has practical advantages. Flexible models interpolate smoothly across parameter space, admit well-behaved likelihoods, and support standard optimisation tools. They allow tensions between datasets to be managed without immediate recourse to model rejection, and they facilitate comparisons between surveys with different systematics and sensitivities. In this sense, parameter freedom functions as a stabilising device, enabling precision cosmology to proceed despite increasing empirical complexity.

The cost of this stabilisation is not primarily technical but epistemic. As flexibility becomes normative, the locus of explanation shifts. Rather than asking why certain cosmological behaviours must occur given a restricted set of assumptions, analysis increasingly asks how much freedom is required to accommodate them. The distinction is not semantic. In the former case, explanation is grounded in necessity; in the latter, it is grounded in fit. The success of a model becomes inseparable from its capacity to absorb variation.

This change also alters the role of falsification. In a tightly constrained framework, failure is informative: it signals that a foundational assumption may be incorrect or incomplete. In a highly flexible framework, failure is deferred. Discrepancies are translated into parameter covariance, degeneracy, or extended confidence regions, and the framework itself remains largely insulated from decisive challenge. This does not render the enterprise unscientific, but it does soften the boundary between explanation and accommodation.

It is important to emphasise that this development has not arisen from negligence or conceptual confusion. It reflects deliberate choices made in response to the scale and sophistication of modern surveys. Forecast-driven optimisation requires models that

can be smoothly varied and efficiently marginalised. Hard constraints, sharp boundaries, and discrete admissibility conditions are methodologically awkward in this setting. As a result, they are rarely foregrounded, even when they might be conceptually relevant.

The consequence is that parameter freedom has come to function not merely as a modelling convenience but as an implicit standard of adequacy. Models that resist deformation are disfavoured, not because they fail observationally, but because they do not integrate easily into the forecasting and optimisation pipeline. This is the context in which concerns about explanatory drift arise. The issue is not that flexibility exists, but that it increasingly defines what it means for a cosmological model to succeed.⁵

3. Instrument optimisation inside a loosened framework

Forecast-driven cosmology is not new, but its reach and influence have grown significantly. Earlier surveys typically tested a small number of tightly specified models, often designed to probe particular assumptions or parameters. By contrast, contemporary forecasts routinely evaluate survey performance across extended parameter families, including flexible dark-energy parameterisations in which the cosmic expansion rate is allowed to vary with time, alongside other phenomenological extensions introduced to absorb departures from the standard model.

Within this framework, success is usually defined by how tightly parameters can be constrained, while complementarity between probes is understood in terms of how efficiently degeneracies can be reduced. These criteria are methodologically effective and technically sophisticated, and they have proved invaluable for comparing surveys and guiding design choices. At the same time, they reflect a subtle but important change in emphasis. Explanatory weight shifts away from asking why a restricted set of models should describe the universe, and toward assessing how well increasingly flexible frameworks can be bounded. That shift is rarely made explicit, but it warrants closer examination.

Large-scale structure observables that probe geometry directly provide a useful test of how forecast practice treats constraint-sensitive information. Cosmic voids are particularly instructive in this regard, because they are often presented as near-geometric probes with minimal dependence on complex baryonic physics. A clear example is the Roman void forecast of Verza et al.,⁶ which evaluates the constraining power of Roman High Latitude Spectroscopic Survey void statistics across Λ CDM and two increasingly flexible dark energy extensions. This is not presented here as a fault of

⁵ Ernan McMullin, *Values in Science, PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* (1982). Classic discussion of methodological values such as explanatory power, simplicity, and accommodation in scientific theory choice.

⁶ Giovanni Verza et al., *Cosmology with Voids from the Nancy Grace Roman Space Telescope*, *Astrophysical Journal* 993, 227 (2025). Forecast study assessing Roman HLSS void statistics under Λ CDM, w CDM, and w_0w_a CDM, used here as a representative example of forecast-driven optimisation within extended parameter spaces.

the authors. It is a representative instance of how forecast practice now defines what questions are treated as central and what counts as success.

The choice of model space is itself revealing. Λ CDM treats dark energy as a cosmological constant, fixed in time. In w CDM, the equation of state parameter w is allowed to depart from minus one, but is assumed constant. In w_0w_a CDM, w is allowed to vary with cosmic time through a two parameter form, so that dark energy can mimic a wider range of expansion histories.⁷ Each step loosens the framework by adding degrees of freedom, increasing its capacity to fit a wider range of behaviours. A forecast then judges success by how tightly those extra freedoms can be bounded, rather than by whether they were needed in the first place.

Modern cosmological surveys are no longer justified primarily by the data they will collect, but by what those data are expected to constrain. Before an instrument is built or a survey begins, its scientific value is assessed through forecasts that estimate how well it will bound selected cosmological parameters. These forecasts are central to decisions about survey design, observing strategy, and investment. They also shape, often implicitly, what kinds of theoretical questions are treated as answerable.

The tools used for this purpose are well established.⁸ Fisher matrices and related techniques estimate how changes in survey design translate into tighter parameter bounds. Figures of merit allow different experiments to be compared on a common scale. Within this framework, a successful survey is one that improves constraints efficiently, reduces parameter degeneracies, and complements existing probes. These are practical and understandable goals. They make complex projects manageable and provide funders and collaborations with clear criteria for success.

At the same time, these methods favour certain kinds of models. Forecasts work best when parameters vary smoothly and continuously, so that small changes in the data can be translated into small changes in inferred values. Extended parameterisations, such as time-varying descriptions of dark energy, fit naturally into this scheme. They can be adjusted incrementally, marginalised over, and compared across surveys. Models that rely on hard geometric restrictions or sharp admissibility conditions are much harder to treat in this way. They do not degrade gracefully under variation, and they offer little guarantee of a positive return once data are taken.

This asymmetry is not a matter of preference or philosophy. It follows directly from how forecasts are used. When survey justification depends on demonstrating robust

⁷ Eric V. Linder, *Exploring the Expansion History of the Universe*, *Physical Review Letters* 90, 091301 (2003). Introduces the w_0w_a parameterisation used widely in forecast studies to model time-varying dark energy.

⁸ Albrecht, A. et al., *Report of the Dark Energy Task Force*, U.S. Department of Energy and National Science Foundation, 2006. Defines the use of Fisher-matrix forecasts, figures of merit, and degeneracy reduction as standard criteria for evaluating and optimising cosmological surveys, and formalises constraint-based performance as the primary measure of scientific return.

performance across a range of scenarios, models that are flexible by construction are easier to defend. As a result, flexibility becomes a practical asset. It allows surveys to promise useful outcomes even if the universe does not behave exactly as expected.

This logic is visible across generational surveys, and it is particularly clear in Roman related forecast work. Roman’s planned combination of large volume, high tracer density, and spectroscopic precision makes it well suited to studies of large-scale structure. Forecasts of its performance typically highlight gains in constraining power over extended parameter spaces, often emphasising improvements relative to earlier surveys or the benefits of combining multiple probes. The success of the mission is thus framed in advance in terms of how much existing freedom can be narrowed.

Void-based observables provide a useful example. Cosmic voids are sensitive to large-scale geometry and are often described as offering information that complements traditional clustering statistics. In principle, this makes them candidates for testing geometric assumptions directly. In practice, however, void forecasts are usually framed in the same terms as other probes. Their value is expressed through their contribution to tightening constraints on flexible parameterisations and reducing degeneracies when combined with other datasets. The question asked is not whether such freedom is needed, but how well it can be bounded.

A complementary illustration can be found in recent ultra-high resolution weak lensing reconstructions enabled by the James Webb Space Telescope. A 2026 analysis of the COSMOS field by Scognamiglio, D *et al.* presents an exceptionally detailed reconstruction of projected mass inferred from shear correlations, revealing filamentary structure, nodes, and voids with unprecedented clarity.⁹ The technical achievement is considerable. Yet the explanatory claims attached to the result are modestly overstated. What is directly measured is the coherent distortion of background galaxy shapes. The translation of that signal into a “map of dark matter” rests on prior commitments about how spacetime curvature, stress-energy, and structure formation are to be interpreted. The striking agreement between the reconstructed map and standard expectations is therefore unsurprising. The reconstruction pipeline and the interpretive framework are co designed. Under an alternative geometry first framework, the same data would be taken as confirmation that global constraints imprint coherent lensing structure across scales, rather than as evidence for a particulate dark sector. The result constrains reconstruction fidelity and internal consistency, but it does not discriminate between ontological accounts. Its persuasive force lies largely in visual and rhetorical confirmation, not in explanatory exclusion. In this respect, ultra-high-resolution weak lensing plays a role analogous to void-based observables at the forecasting stage: both

⁹ D. Scognamiglio, G. Leroy, D. Harvey, R. Massey, J. Rhodes *et al.*, *An ultra-high-resolution map of (dark) matter*, *Nature Astronomy* 2026 (in press). Presents an exceptionally detailed weak-lensing reconstruction of projected mass in the COSMOS field from *JWST* imaging, achieving twice the angular resolution of previous space-based mass maps and revealing coherent filamentary and cluster structure across $0.77 \text{ deg} \times 0.70 \text{ deg}$.

are geometrically sensitive probes whose capacity to discriminate is muted by their systematic embedding within flexible interpretive frameworks.

This way of framing scientific return is effective, and it has delivered clear guidance for survey design. It also has consequences. When instruments are optimised to perform well within flexible frameworks, those frameworks become entrenched through repeated use. Alternative approaches that rely on sharper restrictions are not rejected outright, but they struggle to enter the optimisation process at all. They are difficult to forecast, hard to compare, and risky to justify.

The result is a feedback loop between instrumentation, forecasting practice, and model choice. Surveys are designed to constrain flexible models; those models are favoured because they integrate smoothly with forecast tools; and success is measured by how efficiently remaining freedom is reduced. This does not arise from any single decision or actor. It is the cumulative outcome of rational choices made under uncertainty. Nonetheless, it has the effect of steering cosmological explanation toward accommodation rather than exclusion, even in cases where the data might support sharper questions.

4. How tension is handled in a forecast-driven framework

Disagreement between cosmological measurements is no longer unusual. Differences in inferred expansion rates, clustering strength, or growth history have persisted across multiple surveys and probes, even as observational precision has improved. These tensions are not peripheral anomalies; they arise in precisely those regimes where modern surveys are most powerful and investment is greatest.

Within a forecast-driven framework, such discrepancies are rarely treated as decisive. Instead, they are managed.

Flexible parameterisations allow differences between datasets to be translated into shifts in best-fit values, correlations among parameters, or broader confidence regions. When multiple probes disagree, the usual response is not to question the underlying framework, but to ask whether additional freedom is required to bring the results into statistical alignment. In this way, tension is absorbed rather than resolved. The point is not that this is careless, but that it changes the function that tension plays.

This approach has clear practical advantages. It allows heterogeneous datasets to be combined without immediate fragmentation of the framework, and it avoids premature claims of failure in the presence of complex systematics. From an institutional perspective, it is also reassuring. Surveys justified on the basis of constraining power are unlikely to return empty-handed results if the framework can adapt to accommodate what is observed. Scientific return remains measurable, even in the presence of disagreement.

The effect of this logic, however, is to change what tension signifies. In a tightly constrained framework, persistent disagreement between observations and predictions carries strong diagnostic weight. It forces reconsideration of assumptions about geometry, symmetry, or boundary conditions. In a highly flexible framework, the same disagreement is more easily reclassified as an indication that additional freedom is required. The question shifts from “why does the universe behave this way?” to “how much freedom is needed to accommodate this behaviour?” The framework adapts, and the explanatory demand is softened.

This distinction is not merely rhetorical. It affects how failure is recognised and how theoretical risk is distributed. When flexibility is abundant, the burden of adjustment falls on parameters rather than on structure. The framework itself becomes resilient to challenge, while individual parameter values remain negotiable. As a result, tension no longer functions as a clear signal of explanatory inadequacy but as an impetus for further calibration.

It is important to stress that this mode of response is not irrational. Cosmological data are indirect, systematics are difficult to model exhaustively, and premature rejection of a broadly successful framework would be imprudent. Yet there is a difference between caution and deferral. When discrepancies persist across independent probes and epochs, the repeated expansion of parameter space risks transforming tension from a source of insight into a routine feature of analysis.

The issue, then, is not whether current models can be made to fit the data. In most cases, they can. The deeper question is what is gained, and what is lost, when concordance is maintained primarily through accommodation. Explanation, in its traditional sense, involves showing that observed regularities follow from a restricted set of principles or conditions. Accommodation, by contrast, ensures consistency by widening the range of permissible behaviours. Both have their place, but they serve different epistemic roles.

At present, cosmology appears increasingly comfortable with the latter. Tensions are managed rather than resolved, and the success of a framework is judged by its resilience rather than by its necessity. This does not constitute a crisis, but it does mark a change in explanatory posture. Recognising that change is a prerequisite for assessing whether it is a temporary expedient or a more permanent redefinition of what cosmological explanation is taken to require.

5. Constraint-first approaches as a contrasting explanatory posture

The methodological developments described above are not the only way cosmology has sought explanation. Alongside the dominant tradition of parameter extension, there exists a smaller but well-established body of work in which explanatory power is pursued through the deliberate restriction of admissible spacetime structures. These approaches are united not by shared conclusions, but by a common posture: explanation is achieved by narrowing the space of possibilities rather than by enlarging it.

A canonical example is the Weyl curvature hypothesis proposed by Roger Penrose.¹⁰ Rather than treating low initial entropy as an emergent or statistical outcome, the hypothesis imposes a geometric constraint on the universe's initial state, sharply limiting the allowed configurations of spacetime curvature. The force of the proposal lies not in its empirical success or failure, but in its explanatory ambition. It accepts severe restriction, and with it the risk of falsification, as the price of necessity. The hypothesis does not interpolate; it excludes.

A similar commitment to structural discipline is evident in isotropic singularity frameworks developed by George Ellis and Paul Tod.¹¹ In these models, the initial cosmological boundary is subject to stringent geometric regularity conditions. Generic solutions are ruled out, not accommodated. The aim is not to fit late-time phenomenology but to establish whether a restricted class of spacetimes can account for observed large-scale order without recourse to adjustable parameters. Here again, explanation is inseparable from limitation.

These constraint-first approaches differ widely in motivation and scope, and none offers a comprehensive alternative to contemporary concordance cosmology. That is not their significance in the present context. Their importance lies in demonstrating that cosmology has, within its own mainstream history, treated hard constraints as a legitimate and even necessary route to explanation. They embody a standard in which failure is informative and admissibility is earned rather than assumed.

The Time-Symmetric Parity Constraint is briefly noted here as a recent attempt to pursue explanation through a hard geometric restriction, by imposing a global algebraic condition intended to limit admissible cosmological histories.¹² Its relevance in this essay does not depend on its ultimate correctness, but on the methodological contrast it helps to make visible.

The contrast with contemporary parameter-centric approaches is therefore methodological rather than doctrinal. Constraint-first models prioritise necessity over coverage and accept reduced flexibility as a virtue. Parameter-first models prioritise resilience and continuity, deferring exclusion in favour of accommodation. Both strategies respond to genuine challenges in cosmology, but they embody different answers to the question of what explanation demands.

¹⁰ Roger Penrose, *Singularities and Time-Asymmetry*, in *General Relativity: An Einstein Centenary Survey* (1979). Proposes a geometric constraint on initial spacetime curvature to explain low entropy and temporal asymmetry.

¹¹ George F. R. Ellis and Paul Tod, *Isotropic Singularities and Their Stability*, *Classical and Quantum Gravity* 16, 3647–3668 (1999). Develops cosmological models in which strong geometric regularity conditions are imposed at the initial boundary.

¹² Ian Salisbury, *The Time-Symmetric Parity Constraint*, Paper 3, *ante*. Develops a cosmological framework based on an imposed algebraic constraint on spacetime geometry, proposed as an alternative to parameter-driven extensions.

Making this contrast explicit does not require choosing sides. It does, however, clarify what is at stake. As cosmology continues to refine its observational reach, it must decide whether explanation is best served by ever-greater flexibility or by renewed willingness to say that some possibilities are simply not allowed. The choice is not technical, but methodological, and it will shape how future discoveries are understood.

6. Concluding remarks

The argument developed here is not that contemporary cosmology is misguided, nor that its dominant frameworks are on the brink of collapse. It is, rather, that the standards by which cosmological explanation is judged are undergoing a quiet but consequential shift. Forecast-driven practice, expanded parameter spaces, and optimisation across flexible models have become methodological norms, shaping how data are anticipated, interpreted, and reconciled. These developments have enabled remarkable empirical progress. They have also altered the balance between constraint and accommodation.

At earlier stages in the field, explanatory success was closely tied to structural necessity. Geometric assumptions, symmetry principles, and boundary conditions sharply limited the space of admissible models, and tension functioned as a diagnostic tool. In the present regime, explanation increasingly takes the form of managed concordance, achieved through calibrated flexibility. Discrepancies are absorbed rather than confronted, and frameworks are valued for their resilience as much as for their necessity.

This shift is understandable, and in many respects pragmatic. Modern datasets are complex, systematics are subtle, and premature exclusion would be costly. Yet pragmatism carries its own risks. In a landscape shaped by long-horizon capital investment and advance survey justification, methodological convenience and institutional reassurance can become difficult to disentangle. When flexibility becomes the primary criterion of success, explanation can be thinned to accommodation, and the distinction between understanding and fitting can blur. The concern raised here is not that cosmology has chosen flexibility, but that it may do so without fully acknowledging what is being traded away.

Constraint-first approaches, whether historical or contemporary, remind us that alternative standards remain available. They accept narrow admissibility, explicit failure modes, and heightened risk as integral to explanation. They may ultimately fail, but they fail honestly. Their continued relevance lies less in their conclusions than in the discipline they exemplify.

As cosmology enters an era of unprecedented observational precision, it faces a choice that is methodological rather than technical. It can continue to refine its capacity for accommodation, or it can ask again what must be the case for the universe to be as it is.

Making that choice explicit does not weaken the field. On the contrary, it clarifies what cosmology now seeks to explain, and why.

[Go to the catalogue](#)