Opaque Surprises — Where Scenario Planning Falls Short

ABSTRACT

Scenario planning earned its reputation as one of the twentieth century's most effective tools for strategic decision-making under uncertainty. This essay examines whether the methodology remains fit for purpose in an environment shaped by frontier artificial intelligence, drawing on several decades of practitioner experience in corporate and institutional foresight work.

The argument proceeds by distinguishing three kinds of strategic surprise. Continuous surprise — change within visible trend lines — is the terrain scenario planning was built for and continues to handle well. Discontinuous surprise — rupture whose general character can still be named in advance, such as political upheaval or financial crisis — has long been accommodated, if imperfectly, through dedicated discontinuity scenarios. Opaque surprise is different in kind: a discontinuity whose character cannot be specified in advance, because the technology producing it surprises even the people building it. Current AI development falls into this third category. No refinement of narrative scenario methodology can compensate for it.

Rather than abandoning scenario planning, the essay proposes a more limited and more precise posture for opaque surprise. That posture rests on three things: robust strategy over optimal strategy, shorter review cycles, and explicit acknowledgment when a branch of the future cannot yet be narrated. It closes by suggesting that trilemma models — which structure the constraints on a system rather than narrate its content — may offer more durable guidance under these conditions than scenario narratives can.

I.

Back in the day, in what is now ancient history, I was part of a corporate strategy team that developed five scenarios for what the telecommunications industry might look like a decade out. Three were continuity scenarios, distinguished by which layer of the industry would drive its evolution — one built around network infrastructure, one around services, one around devices. The other two were discontinuity scenarios: one in which the industry was reshaped by a major technological breakthrough, the other by political upheaval.

We presented all five to the company's executive management, who were skeptical but ultimately approved them for internal strategic planning and for use in collaborations with clients. There was one significant caveat: the discontinuity scenarios were judged too disruptive for outside audiences and were dropped from external use. Only the three continuity scenarios traveled beyond the company's walls. They were disseminated widely and served Ericsson well through several planning cycles — used internally, at industry forums, and at client workshops held around the world. Participants in the original exercise and the external audiences who encountered the scenarios later differed widely in which future they found most plausible. What nearly everyone agreed on was the sense of urgency the scenarios conveyed. The emerging internet had a way of focusing minds.

That workshop is one data point among many from a long career spent building these tools and watching organizations use, misuse, and occasionally be transformed by them. Scenario planning was — by most reasonable accounting — one of the twentieth century's most valuable strategic instruments. It helped Shell navigate the 1970s oil shocks when its competitors were blindsided. It gave generations of executives a structured way to hold multiple futures in mind at once, rather than betting everything on a single forecast. It earned its reputation through results, in environments at least as turbulent as the one in front of us now.

The question this essay sets out to answer is simpler to ask than to answer: is it still fit for purpose? Not whether it remains useful — it plainly does, for a great deal of what organizations face. The harder question is whether there is now a category of strategic uncertainty the methodology was never built to handle. And whether continuing to apply it there anyway does more harm than the comfort of a plan is worth.

II.
What Scenario Planning Actually Is — and Isn't

Scenario planning occupies an awkward position in the strategic management literature. It is simultaneously over-theorized and under-described — exhaustively codified in methodology texts yet poorly conveyed as a practice. The gap matters. How the methodology functions in institutional settings differs enough from how it is specified in frameworks that conflating the two produces a distorted picture of both its value and its limits.

The basic architecture is well established. Practitioners identify driving forces — demographic, technological, economic, political, environmental — operating on a given system over a chosen planning horizon. From these, they isolate critical uncertainties: forces whose direction or magnitude is genuinely unknown and whose variation would produce qualitatively different futures. Two or three such uncertainties, placed on intersecting axes, generate a matrix of scenario spaces. Each space is then developed into a coherent narrative: internally consistent, causally plausible, and distinct enough from the others to force genuine strategic thinking rather than marginal adjustment.

The distinction between scenario planning and forecasting is foundational and persistently misunderstood, including by practitioners who should know better. Forecasting attempts to specify what will happen, typically by extrapolating trends or aggregating expert judgment. Scenario planning maps what could happen across a range of plausible futures, and asks what strategic posture remains robust across that range. The methodology is not a sophisticated forecasting tool. It is a structured discipline for operating under genuine uncertainty. That is why conflating the two — as executives under institutional pressure frequently do — is not merely an intellectual error but a practical one. An organization that treats its scenarios as predictions has eliminated the methodology's primary value.

A second case, two decades after the Ericsson workshop, made the same point from a different angle. As part of the core team building five-year telecommunications scenarios during my time at Gartner, I watched the scenarios land well — genuinely well, appreciated by carriers and equipment vendors alike, used in planning cycles across multiple client organizations. What I had not expected was how persistently those same clients wanted to convert the scenarios into something they were never built to be. Telcos and vendors kept asking for probabilities: which scenario was most likely, weighted how, by when. The scenarios were sound. The instinct to interrogate them as forecasts rather than as a mapped possibility space was the same instinct named above: mistaking scenario planning for prediction. It was simply playing out, this time, among sophisticated industry clients who, if anyone should have known better, should have.

In practice, the majority of scenario work addresses what might be called continuity futures — worlds in which present trends develop along different but unbroken trajectories. The global telecom restructuring of the late 1990s, for instance, was visible in outline years before it arrived. The critical uncertainties concerned timing, regulatory sequencing, and which incumbent players would prove capable of strategic adaptation — not whether restructuring would occur. Scenario work in that context helped organizations calibrate responses to a range of plausible timelines and competitive configurations, not anticipate a fundamentally different kind of world. This is scenario planning functioning as designed, and functioning well.

Discontinuity scenarios occupy a different register. These address ruptures — events or transitions that break trend lines rather than extend them. Political upheavals, pandemic events, financial system failures, breakthrough technologies that alter competitive structure. Traditional scenario methodology accommodates discontinuities as a distinct branch of the scenario space, typically held somewhat apart from the continuity scenarios precisely because their narrative development requires different assumptions. In practice, however, discontinuity scenarios were always the least methodologically satisfying component of the work. They tended toward the generic: a major conflict, a systemic financial shock, a disease event. The framework offered tools for asserting that discontinuity was possible, and for sketching broad organizational responses, but limited tools for specifying the character of a rupture in advance. The pandemic scenario was a staple of responsible foresight work for two decades before 2020. What scenario work could not specify was the precise epidemiological profile, the specific institutional failures, the particular pattern of economic disruption, or the political dynamics that would follow. The scenario named the category. It could not map the territory.

This distinction — between naming a class of discontinuity and specifying its character — is the crack in the methodological foundation that current conditions are widening. Understanding why requires first examining what acceleration does to the planning environment. It then requires examining what a particular kind of technological change does that acceleration alone does not.

III.
The Acceleration Problem

A rule of thumb common among scenario practitioners held that calibrating ten years forward required examining twenty years back. The logic was corrective rather than mystical: doubling the look-back window disciplined analysts against the natural tendency to underestimate how much actually changes in a decade. It's a bias that flattens even careful forecasting. Applied carefully, the rule produced useful humility. It forced the question of how strange the present would look to an observer transported forward from two decades earlier. That strangeness became a calibration device for the future.

The rule rests on an assumption worth making explicit, because it is precisely the assumption current conditions strain. It assumes the rate of change is increasing. But it also assumes that the nature of change remains legible across the doubled timeframe: that today's world would strike a visitor from two decades ago as remarkable, yet a sufficiently careful analyst working a decade ago could, with effort, have identified most of the forces that would matter today. The diffusion of smartphones and mobile internet, the build-out of cloud computing infrastructure, the platforms that restructured media, retail, and advertising around social networks and recommendation algorithms — each was visible in outline well before its full effects arrived, even if precise timing and combination were not. The twenty-year rule worked because acceleration changed the speed of arrival, not the discoverability of what was arriving.

In practice, acceleration does real work on the scenario planning process, and it is worth being precise about what that work consists of. It shortens the useful planning horizon, since driving forces that once held for a decade now shift meaningfully within three or four years. It increases the number of critical uncertainties competing for attention at any one time, because faster-moving systems generate more inflection points per unit of calendar time. And it makes scenario narratives age faster than the institutional processes built around them. A scenario set commissioned for a five-year strategic cycle may already look visibly dated by year two, while the organization continues operating as though the narratives still describe live possibilities.

None of this breaks the methodology. Faster change is still legible change. The appropriate response to acceleration is adjustment, not abandonment. Scenarios can be refreshed on shorter cycles, horizon work can be compressed from ten years to five or three, and review rhythms can be matched to the velocity of the domain rather than to an inherited planning calendar. Scenario planning has absorbed acceleration before. The acceleration problem is real, and it is also, in the end, a problem of degree.

What follows is not a problem of degree. It is a problem of kind — and it begins with a closer look at what discontinuity scenarios were actually capable of, even at their best.

IV.
The Limits of the Discontinuity

Section II closed by naming a limitation: discontinuity scenarios could assert that rupture was possible without specifying its character. It is worth examining why that limitation existed, because

the explanation reveals exactly what current AI development breaks that ordinary discontinuity did not.

Even at their most disruptive, traditional discontinuities had a recognizable anatomy. A political upheaval, however sudden, takes place among actors whose interests, capabilities, and grievances are visible to anyone studying the system closely — who holds power, who contests it, which institutions will bend and which will break. The Soviet collapse surprised most observers in its timing and speed, but not in its cast of characters or institutional logic. The relevant forces — economic stagnation, nationalist pressure, a leadership unable to reconcile reform with control — had been cataloged by analysts for years beforehand. A financial crisis draws on a similarly known repertoire: leverage, correlated risk, liquidity mismatch, panic. Each crisis is novel in its trigger and propagation, not in the grammar it speaks.

Technology breakthroughs followed the same pattern, even when they arrived as genuine surprises in scale or timing. The internet's commercial explosion in the mid-1990s is the clearest case from direct practitioner experience. Few forecasters anticipated how quickly it would restructure media, retail, and telecommunications, or how completely it would outrun the regulatory frameworks built to govern those industries. But the ingredients were not hidden. Packet-switched networking, falling computing costs, the slow accumulation of protocols and standards, the venture infrastructure assembling around Silicon Valley: analysts working in foresight practice through the early nineties could see these elements gathering. They simply could not specify exactly when or how the elements would combine. The surprise was real. It was a surprise of timing and magnitude, not of fundamental nature. Once the explosion happened, it was explicable in terms already available before it happened.

This is what made the discontinuity branch of scenario work tolerable, even in its weaker moments. A scenario team could not specify when a pandemic would arrive. It could specify — correctly, as it turned out in 2020 — that a fast-moving respiratory pathogen would expose supply chain fragility, strain healthcare capacity, and force abrupt behavioral change at societal scale. The category was nameable even when the instance was not. This is the quiet assumption underneath nearly all discontinuity scenario work: that rupture, however violent, draws its material from a landscape careful observers have already mapped.

It is this assumption — not the pace of change, but the knowability of change's character — that current developments in artificial intelligence place under genuine strain. The question is no longer how quickly the future will arrive, which acceleration alone can answer. It is whether the future, once it arrives, will be describable in terms available beforehand at all.

V.
Continuous, Discontinuous and Opaque Surprise

The distinction worth drawing, then, is not simply between continuity and discontinuity. It is among three qualitatively different kinds of surprise — continuous, discontinuous, and opaque — offered

here as a refinement of standard scenario architecture, not a replacement for it. The first two terms are deliberately continuous with the vocabulary established in Section II; the third names something that vocabulary cannot yet capture.

Before developing these three categories, it is worth pausing at a framework that maps roughly the same territory. It is worth being precise, too, about where that mapping holds and where it breaks down. In February 2002, in a context entirely removed from scenario planning, Donald Rumsfeld offered a formulation that has since entered general use: there are known knowns, things we know we know; known unknowns, things we know we do not know; and unknown unknowns, the ones we do not know we do not know. The political circumstances in which those words were spoken need not detain us here. The epistemological content stands independently of them, and has been taken seriously in domains from intelligence analysis to evolutionary biology to project risk management.

The Rumsfeld categories map approximately onto the three kinds of surprise developed here. Known knowns correspond roughly to the terrain of continuous surprise — the trends we track and understand. Known unknowns correspond to discontinuous surprise — the ruptures we know are possible, whose character we can name even when we cannot time their arrival.

Unknown unknowns come closest to opaque surprise. But the fit is approximate, not exact, and the difference is precisely where the present argument's force lies. Rumsfeld's unknown unknowns are things that exist but have not yet been discovered — gaps in our map of a territory that is itself stable and in principle mappable. The opaque surprise described here is something more specific and more troubling: a discontinuity whose character cannot be specified even by those closest to producing it. The capability emerges from scale in ways that resist prediction from the ingredients themselves. That is a different kind of not-knowing, and it requires a different response.

Continuous surprise describes change within continuity scenarios: developments whose direction is visible well in advance, with genuine uncertainty confined to timing and magnitude. This is the terrain scenario planning was built for and continues to handle well. The telecom restructuring described earlier is paradigmatic — visible in outline years before its arrival, uncertain chiefly in sequencing and which incumbents would adapt. Most strategic uncertainty organizations face, even today, falls into this category. That alone is reason enough not to discard the methodology that addresses it.

Discontinuous surprise is the traditional discontinuity scenario, named with more precision than the literature usually affords it. A rupture breaks the continuity of trend lines, but its character — the grammar it speaks — can be articulated in advance, even if its timing cannot. Political upheaval, pandemic disease, financial crisis: each draws on a recognizable repertoire, and each can be prepared for through explicit discontinuity branches and contingency planning, however imperfectly. Discontinuous surprise is surprising in its occurrence. It is not surprising in its nature.

Opaque surprise is different in kind. It names a discontinuity whose character cannot be specified in advance — not because observers have failed to look carefully enough, but because the nature of the rupture is genuinely unknowable until it occurs. This describes, with some precision, the position researchers building the most capable AI systems currently occupy. Specific emergent capabilities have repeatedly surprised the labs producing them, sometimes only after a model has already been deployed. These are skills and behaviors that appear only past certain scales of training and compute, neither designed nor predicted by the people doing the designing. This is not a failure of planning, foresight, or institutional competence. It appears to be a structural feature of how capability emerges from current AI architectures: a property of the technology, not a gap in anyone's diligence.

The distinction echoes, without quite reproducing, Frank Knight's century-old separation of risk from uncertainty — risk being randomness whose probability distribution is knowable, uncertainty being randomness that resists even probabilistic specification. Discontinuous surprise fits comfortably within Knightian uncertainty: the probability of a given discontinuity is unquantifiable, but the space of possible discontinuities is enumerable in advance. Opaque surprise pushes past even Knight's category. It is not simply that the probability of a rupture is unknown — the contents of the possibility space itself are unknown until the rupture has already occurred. It is uncertainty about uncertainty: a meta-level opacity that classical decision theory was not built to formalize. Scenario planning, built on enumerable possibility spaces, was not built to narrate it either.

The implication runs further than it first appears. Opaque surprise does not merely compress the planning horizon, the way acceleration does. It challenges the foundational premise of scenario methodology itself — that careful observation of present ingredients allows a skilled analyst to anticipate the character of future change, even when timing remains uncertain. The internet's builders could be surprised by how fast their creation spread without being surprised by what kind of thing it was. The researchers training today's frontier AI systems are reporting something different: surprise at what kind of thing they have built. That surprise arrives after the fact, discovered by interrogating systems whose relevant capabilities were not specified by their designers, and were not visible by inspecting the ingredients beforehand. When the people with the closest possible vantage point on a technology cannot anticipate its character eighteen months out, the scenario planner working at one remove has no comparative advantage left to draw on — however skilled, however rigorous in driving-force analysis. The toolkit that worked for discontinuous surprise simply does not transfer.

VI.
What This Means for Practice

None of this argues for abandoning scenario planning, and it would be a mistake to read it that way. Continuous and discontinuous surprise still account for the large majority of the strategic uncertainty most organizations actually face — markets shift, regulations change, competitors rise

and fall. Even genuine discontinuities mostly retain a nameable character. The methodology remains valuable for exactly the territory it was built to cover. What it needs is not abandonment but clarity: an explicit acknowledgment, built into the practice itself, of where its writ runs out.

For the portion of strategic uncertainty that is genuinely opaque, a different posture serves better than refined narrative scenarios. Robust strategy — organizations and policies built to absorb surprises whose character is unknown — does more useful work here than optimal strategy does. Optimal strategy is calibrated against named futures that may simply be wrong about what kind of future is coming. This is a familiar idea in risk management generally. What is new is recognizing that it now applies to a portion of technological strategy that scenario teams have historically tried to narrate rather than merely buffer against.

Cycle time matters more than horizon length under these conditions. Ten-year scenario work, however well executed, asks an organization to commit to a narrative architecture for a period during which the underlying technology may pass through several genuinely unanticipated capability thresholds. Shorter cycles better match the rate at which the relevant uncertainty actually resolves, or fails to: annual or even quarterly review rhythms for the opaque-surprise component of a strategy, layered alongside longer-cycle work for the continuous-surprise component.

Perhaps the most important practical change is cultural rather than methodological. Scenario workshops should be willing to say, of one branch of the future, that it cannot currently be narrated — and to treat that statement as a finding, not an admission of failure. A workshop that produces four polished, internally consistent AI scenarios may be more dangerous than one that produces three narratable scenarios and an acknowledged gap. The polish disguises a confidence the underlying methodology cannot actually support.

The trilemma model used in earlier development scenario work retains its value here for a specific structural reason. It is worth pausing to define the term before explaining why. A trilemma model specifies three goals or values that cannot all be fully achieved at once. Any real system has to trade off some combination of the three, and different configurations represent different compromises rather than a single correct answer. The classic illustration comes from international economics: the so-called impossible trinity holds that a country cannot simultaneously maintain a fixed exchange rate, free movement of capital across its borders, and an independent monetary policy. It can have any two. Pursuing all three at once is not merely difficult — it is structurally unstable, as a series of currency crises have demonstrated when governments tried anyway. The energy trilemma, widely used in energy-transition planning, applies the same logic to a different domain: it frames national energy strategy as a standing balance among security of supply, affordability of access, and environmental sustainability. No country achieves all three at once. Energy policy is, in practice, a continuous negotiation over which two to prioritize and how much strain the third is allowed to absorb.

What makes this structure useful for opaque surprise is that it frames futures around the constraints a system must reconcile — competing pressures any outcome will have to balance. It does not frame

them around the narrative content of what that outcome will look like. A trilemma can be fully specified without knowing which AI capability will emerge next, because it describes the shape of the trade-off rather than the substance of the surprise. That makes it more durable under opaque surprise than narrative scenarios are. It is also worth more attention than it has typically received outside the development-planning contexts where it originated.

VII.
Closing

None of this is offered as a requiem. Scenario planning does not need to be mourned. It needs to be practiced with more precision about what it can and cannot do — a different and more useful project than declaring the methodology obsolete just because the technology makes for an unusually dramatic obituary.

The most useful thing a scenario planner can offer an organization today may not be a set of narratives at all. It may be the discipline of structured uncertainty: the ability to say clearly which parts of the future are nameable and which are not, and to build strategy that does not pretend otherwise. Knowing what you don't know, and being precise about why you don't know it, is itself a form of strategic intelligence. It may be a more durable one, in fact, than any individual scenario set has ever been. Organizations have always paid scenario planners for stories. The more valuable service, going forward, may be the disciplined refusal to tell one where none can yet be told.

This essay has tried to make that case by tracing a single thread: from the twenty-year rule, through the quiet limitations of discontinuity scenarios that practitioners have long sensed but rarely stated plainly, to a sharper distinction among continuous, discontinuous, and opaque surprise. It is a first attempt, not a settled framework, and it leaves at least two questions open. How should organizations actually make decisions — not workshop exercises, but live capital and personnel decisions — when one branch of the relevant future is genuinely opaque rather than merely uncertain? And does the trilemma model, built originally for development planning in conditions far removed from frontier AI research, actually transfer to this context, or does it merely offer the comfort of a familiar structure applied to an unfamiliar problem? Both deserve more space than a closing paragraph can give them. They are, perhaps, where this inquiry goes next.

Mike Edholm has spent three decades advising corporations and institutions on scenario planning and strategic foresight, with assignments spanning telecommunications, energy, and international development. This is the first essay in a series on scenario planning methodology.