BREAKTHROUGH INSTITUTE

STAKEHOLDER STRATIFICATION

RISK FRAGMENTATION AND DISTRIBUTED DECISION SYSTEMS IN INFRASTRUCTURE DEPLOYMENT

ADAM STEIN

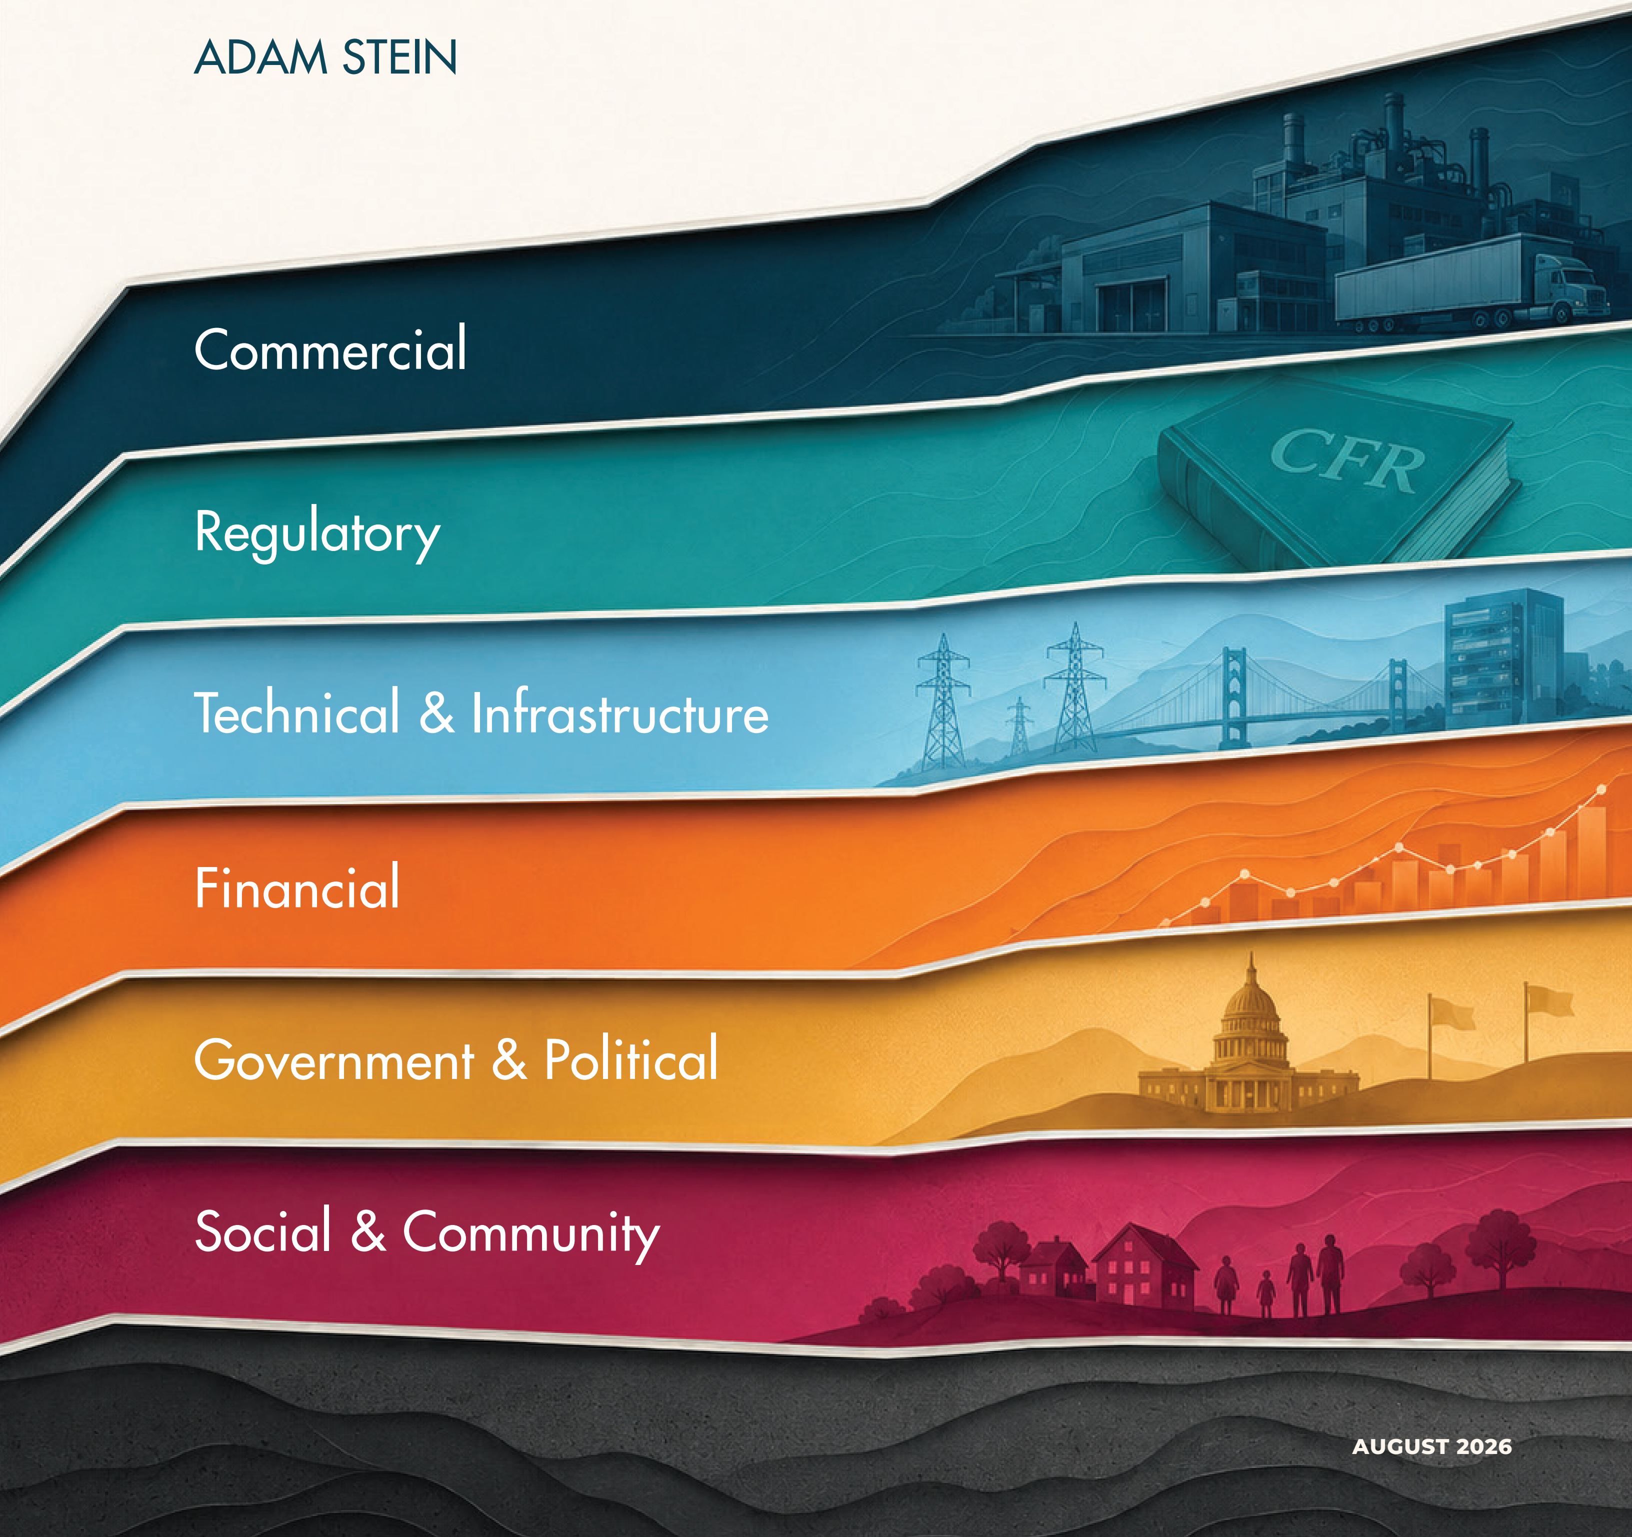

AUGUST 2026

Executive Summary

Large infrastructure projects are governed by distributed decision systems: configurations in which multiple independent stakeholders—each with distinct mandates, evaluation frameworks, and definitions of acceptable risk—must simultaneously reach acceptable outcomes within their own domains. No single actor controls the result. No single intervention is sufficient to produce success. Analysts routinely document projects that exceed cost and schedule estimates, that collapse after years of apparent progress, or that succeed technically but fail commercially or politically. Project failure patterns are not explained by individual factors alone, which are primarily symptoms of the structure in which those factors operate.

This report corrects several specific misunderstandings and misframings that have been adopted broadly. Authority in large infrastructure is not centralized in any one institution. Progress in one domain does not advance the system unless other domains converge. Partial progress is the expected pattern, not an anomaly.

Progress, and failure, are the result of a more interconnected and interdependent structure, and need a different framework. Project realization depends on non-fungible approvals—decisions across legal, commercial, financial, technical, political, and social domains that do not substitute for one another—and outcomes are determined by joint satisfiability, the requirement that all approval conditions hold simultaneously.

In parallel, feasibility is bounded by constraints that no single actor controls. Some are genuinely external (i.e., exogenous), such as macroeconomic conditions. Others are quasi-exogenous: supply chain capacity and infrastructure limits produced by the system's own institutional structure and demand uncertainty, yet experienced by each participant as external conditions they did not create and cannot individually resolve. Still others are internal, manufactured by the system's own risk-allocation decisions. Even when alignment is achieved, execution depends on conditions that can change after approvals have been secured.

Importantly, the concept of "risk" is an institution-specific construct: stakeholders may not disagree about the magnitude of a common risk, but evaluate different categories of uncertainty entirely.

The implications are structural. Single-domain interventions displace constraints rather than resolving them. Stakeholder conflict reflects different risk definitions, not different facts. Demand-side acceleration exposes bottlenecks more quickly; it does not remove them. Effective action requires understanding not only where a constraint appears, but how it is connected to other domains and where it will re-emerge when displaced.

Nuclear energy projects are the primary case in this report because the distributed decision dynamics are most visible there: the number of independent approval domains is large, the time horizons are long, and the consequences of domain-level failure are well-documented. However, the mechanisms are not nuclear-specific.

This report does not attempt to quantify the probability of any specific project's success, prescribe particular interventions, or rank alternative approaches. It develops a structural framework intended to support more accurate analysis, more robust project design, and more disciplined evaluation of policy and institutional interventions.

Section 1 — Introduction: The Structure Behind the Pattern

A consistent empirical pattern runs through large infrastructure development: projects routinely exceed their cost and schedule estimates by wide margins, frequently stall after years of apparent progress, and sometimes collapse after major capital commitments have been made. The explanations offered for this pattern are familiar. They center on regulatory delay, financing cost, political opposition, supply chain constraints, and—most influentially in the academic literature—cognitive and strategic failures by the agents who plan and promote these projects.

Bent Flyvbjerg's work on large-project performance is the most systematic treatment of this pattern. Across a dataset of hundreds of projects spanning decades and infrastructure types, Flyvbjerg documents consistent, significant cost overruns and schedule extensions. His explanation for this pattern—optimism bias and strategic misrepresentation, or what he calls "delusion and deception"—is the dominant account in the field. It attributes the pattern to agent-level failures: planners who underestimate costs through cognitive bias, promoters who misrepresent projections to secure approval.

The account is correct within its scope, but the scope is limited.

Flyvbjerg's analysis is conducted one project, one estimate, one agent at a time. Examined at that resolution, the only causes in view are agent-level, and the explanation is necessarily psychological. But an agent at any single stratum of a large project has bounded authority: even one who correctly diagnoses what would resolve the whole system cannot act outside their domain. The failure pattern therefore persists even when bias is fully removed, because it requires only that authority be distributed, not that anyone be biased or dishonest. Bias is, at most, a contributing factor.[1] It is not the cause.

Flyvbjerg's dataset compounds this problem structurally. It consists of completed projects—projects that were built. Failed, abandoned, or screened-out projects, by definition, are absent from his sample. This dataset truncation is an instantiation of both survivorship bias and selection bias. Yet projects that collapse when a single approval domain becomes binding are precisely the cases a structural account most needs to explain, and precisely the cases missing from the sample. The documented pattern of cost overruns is real. The agent-level explanation for it is incomplete, because the analysis cannot see the cases where the system failed earlier or more completely.

This report widens the frame. It locates the driver in the decision architecture rather than in the agents who operate within it. The relevant question is not whether individual actors are biased or honest, but whether the structure in which they operate is capable of carrying a project to completion across independent approval domains under evolving constraints.

Nuclear energy projects make this structure most visible. They sit at the extreme of nearly every dimension that shapes how complex projects are decided: the number of independent approvals required, the length of the time horizon, the formalization of risk evaluation, and the stringency of the standards applied—which routinely exceed those imposed on demonstrably more hazardous activities and projects such as coal-fired power generation or transportation. This extremity is not evidence that nuclear is a special case. It is the same structure common to large, complex, multi-stakeholder projects, expressed at its limit. Dynamics that are faint in a transmission line or a port expansion are loud and measurable in a reactor project. The framework developed here is presented with that broader applicability in mind.

What appears as a sequence of approvals is more accurately a distributed decision system, in which multiple independent stakeholders must simultaneously reach acceptable outcomes within their own domains. Projects fail because multiple independent conditions must all be satisfied at once, not a sole dominating factor at a fixed point in time. A project may satisfy regulatory requirements but fail to secure financing. It may attract capital but lack a buyer. Even when both conditions are met, it may fail due to supply chain or infrastructure constraints that were not binding when the approvals were granted but became binding during execution. These outcomes reflect the structure of the system. They are expected, not anomalous.

This distributed decision system operates alongside a distinct, parallel system of binding constraints, which can alter decisions and determine whether an approved project can ultimately be executed. Macroeconomic conditions affect the cost of capital. This is an external constraint, which is imported into the system rather than being created by the stakeholders within it. Supply chains determine whether projects can be physically delivered. These constraints are not independent of the stakeholder system; they are partially shaped by regulatory requirements, demand uncertainty, and risk allocation decisions. They are therefore quasi-exogenous to the same distributed decision structure that governs approvals, even when they appear external at the point of decision. Other constraints are hybrid and straddle the line between purely external and quasi-exogenous.

Three conditions recur throughout these systems. **Joint satisfiability** is a property of the approval conditions: the requirement that all of the non-fungible approvals hold at the same time, since none substitutes for another. **Joint acceptability** is a property of the decision-makers: because each approval is held by an independent actor applying its own definition of acceptable risk, meeting that requirement depends on every actor judging the project acceptable in its own domain simultaneously—satisfiability is the structural requirement, and acceptability is how it is achieved, or missed, among independent deciders. **Joint feasibility** is a property of execution: even a fully approved project must be physically, economically, and institutionally deliverable under prevailing constraints, and must remain so as those constraints change. Approval turns on the first two, execution on the third, and realization requires all three at once.

This report analyzes such properties and interactions between the two systems: 1) a distributed decision system, in which approvals are independent and non-substitutable, and 2) a set of binding constraints (external, quasi-exogenous, and hybrid). The objective is not to identify a dominant bottleneck. It is to clarify how decisions are made, how coordination fails, and how constraints emerge from the structure of the system rather than from the failures of the agents operating within it.

Section 2 — Institutional Context: How Infrastructure Projects Are Actually Authorized

Standard analyses of infrastructure development begin with formal processes—licensing pathways, permitting requirements, procurement mechanisms. These processes define points of authorization, but they do not describe how authorization actually operates. They imply a structured progression through identifiable steps. What they obscure is that authorization is not governed by a single process but by the interaction of multiple independent institutional domains, each with its own mandate, timeline, and definition of acceptable outcomes.

These domains do not coordinate through a central authority. Each must independently reach an acceptable outcome, and failure in any one prevents the project from proceeding regardless of success in the others.

At a minimum, large complex infrastructure projects require simultaneous satisfaction across several distinct approval channels. The channels described below are the ones that recur across large projects; the set is representative rather than fixed, and any one of them can stratify into finer, independent decisions. Nuclear projects navigate this structure in its most elaborated form.

Legal and Regulatory Approval. Federal and state licensing, along with environmental permitting, establishes whether the project meets statutory requirements for safety, environmental impact, and site suitability. For nuclear, this channel is administered primarily by the Nuclear Regulatory Commission. The equivalent authority varies across infrastructure types—the Federal Energy Regulatory Commission for transmission, the Army Corps of Engineers for waterway infrastructure—but the structural role is identical: a gatekeeper with statutory authority whose approval is necessary but not sufficient.

Commercial Approval. A viable project requires a committed buyer. Without a long-term revenue commitment, the project has no economic basis. This approval is not granted by a regulatory body; it emerges from a commercial negotiation. Its absence is not a rejection but an absence—no buyer, no project—and it cannot be substituted by success in any other domain.

Financial Approval. Capital providers evaluate whether the project meets investment or lending criteria: construction risk, revenue certainty, macroeconomic exposure, and expected return. Their decision is independent of regulatory approval and independent of commercial commitment. A project can be fully licensed and commercially contracted and still fail to attract capital if financial criteria are not met.

Technical and Infrastructure Approval. Integration into a system—through interconnection, system reliability review, and access to systems including transmission or water—requires the independent authorization of grid operators and reliability organizations. A project cannot generate revenue from power it cannot deliver, and it cannot deliver power it cannot connect to the grid.

Political and Governmental Approval. Policy decisions, public funding conditions, and administrative posture determine feasibility conditions even where formal approval is not required. Political support is often invisible when present and decisive when absent.

Social and Local Approval. Communities and public stakeholders shape outcomes through formal participation processes and informal pressures that propagate into regulatory and political decisions. Social approval is not granted by any single authority; it is maintained continuously and can be withdrawn at any stage.

Each channel is itself internally stratified. For example the Commercial channel: a buyer willing to purchase, a seller willing to supply, and a builder willing to construct at acceptable risk are distinct decision-makers, each holding an independent veto. A buyer may want a product no seller will offer on acceptable terms; a buyer and seller may agree while no builder will accept the construction risk or has the capacity to deliver on the required timeline.

The specific channels vary across infrastructure types—a transmission line faces a different regulatory structure than a reactor, and an LNG terminal faces a different one still—but the underlying structure does not. Every large, complex infrastructure project requires simultaneous satisfaction of legal, commercial, financial, technical, political, and social conditions. The content of each domain differs; its necessity does not.

This structure is often interpreted as sequential—licensing first, then financing, then construction. But, decisions are made in parallel, under uncertainty, with each stakeholder making contingent assessments based on what other domains are expected to produce. Financial institutions assess projects based on anticipated regulatory outcomes. Buyers evaluate commitments based on projected cost and delivery timelines. Developers advance licensing without certainty about financing or demand. These parallel conditional decisions, made without centralized coordination, constitute the system's coordination problem. No one is in charge of the whole, and the whole depends on everyone.

Historically, regulated utilities partially internalized this coordination function. As load-serving entities responsible for planning, procurement, and system integration, they aligned multiple approval channels within a single organizational structure. This reduced coordination complexity but did not eliminate the underlying constraints. The domains remained independent; utilities simply operated within more of them simultaneously.

That model is now dissolving at its edges. Rapid growth in electricity demand—driven primarily by data center development and artificial intelligence infrastructure—has introduced a new institutional class of market participants. Hyperscale technology firms are not utilities. They do not plan, build, or operate power systems. But they have entered the energy market with financial scale and strategic urgency sufficient to engage directly in generation procurement, project finance, interconnection requests, and market rule proceedings. Their presence redistributes coordination across domains; it does not resolve the coordination problem. Commercial approval may be accelerated. Regulatory, technical, and supply chain constraints remain binding and have in some cases become more visible precisely because commercial alignment arrived faster than the other domains could follow.

Not all constraints take the form of institutional approval. Supply chain capacity, workforce availability, and macroeconomic conditions impose limits that operate alongside the approval system. These constraints interact with institutional decisions—regulatory requirements influence supply chain participation; demand uncertainty affects supplier investment—and must therefore be understood as part of the same system, not as conditions that can be addressed after approvals are secured.

Section 3 — Stakeholder Stratification: Multiple Valid Classification Systems

Understanding which stakeholders are involved in a large infrastructure project is not the same as understanding how those stakeholders determine its outcome. The classification system used to organize stakeholders determines what appears to be the relevant constraint, where interventions seem appropriate, and what coordination looks possible. Different classification frameworks produce different apparent structures.

This matters analytically because it explains a persistent failure mode: stakeholders in one domain misread the behavior of stakeholders in another because each is applying a different organizing framework to the same situation. A developer applying a functional-role framework sees a regulator as a gatekeeper whose requirements must be satisfied. The regulator, applying a safety-compliance framework, sees the developer as a risk source whose claims must be verified. Both descriptions are accurate. They are not compatible. The coordination problem is partly a classification problem.

This section presents four classification frameworks—institutional role, sector type, functional role, and decision logic—and examines how they interact. The examples are drawn from nuclear deployment, but the classification structures apply to any large infrastructure project requiring distributed authorizations.

3.1 Classification by Institutional Role

The most direct classification groups stakeholders by their institutional function and formal authority in the project.

Table 1. Institutional Layer Classification of Stakeholders

INSTITUTIONAL LAYER	PRIMARY FUNCTION	REPRESENTATIVE STAKEHOLDERS
Political / Sovereign	Policy direction, public funding, political support	Congress, state governments, executive agencies
Regulatory	Safety licensing and oversight	NRC, sector-specific regulators
Environmental / Permitting	Environmental review, land use approval	Federal/state/local permitting bodies

INSTITUTIONAL LAYER	PRIMARY FUNCTION	REPRESENTATIVE STAKEHOLDERS
Grid / Infrastructure	System integration and reliability	ISOs/RTOs, transmission owners
Commercial (Buyers)	Revenue commitment	Utilities, hyperscalers, industrial buyers
Financial	Capital provision and risk allocation	Banks, equity investors, public finance
Supply Chain / Delivery	Design, manufacturing, construction	Vendors, EPC firms, suppliers
Public / Civic	Legitimacy, local acceptance	Communities, NGOs, media

This framework reflects how projects are formally structured—who holds what authority, what approval is required from whom. It is the framework most visible in project management and regulatory documentation. Its limitation is that it treats institutions as isolated actors rather than as participants in a decision system with cross-domain interactions. It shows the structure but not the dynamics.

3.2 Classification by Sector Type

A second classification groups stakeholders by ownership structure, which determines accountability, decision incentives, and risk tolerance.

Table 2. Stakeholder Classification by Sector Type

SECTOR TYPE	DEFINING CONSTRAINTS	REPRESENTATIVE STAKEHOLDERS
Public	Statutory mandate, procedural accountability, political oversight	Regulators, policymakers
Private	Capital allocation, return thresholds, risk exposure	Banks, developers, hyperscalers
Quasi-Public	Hybrid regulatory and commercial obligations	Utilities, ISOs/RTOs
Civic / Public Sphere	Legitimacy, trust, public influence	Communities, NGOs, media

While public agents' constraints are procedural and statutory, private agents' primary constraints are return thresholds and risk exposure. Hybrid structures operate under both sets of constraints simultaneously, which often produces conservative behavior at the system level even when neither the regulatory nor the commercial logic alone would produce it. This distinction explains why and how each stakeholder responds to external shocks, policy changes, and risk reallocation.

3.3 Classification by Functional Role

A third classification organizes stakeholders by their functional role in enabling or preventing project realization—a framework that reveals where effective control resides, including in places where formal authority is absent.

Table 3. Stakeholder Classification by Functional Roles

FUNCTIONAL ROLE	PRIMARY CONTROL MECHANISM	FAILURE MODE IF ABSENT	REPRESENTATIVE STAKEHOLDERS
Gatekeepers	Legal and regulatory approval	Project cannot proceed legally	NRC, permitting agencies
Buyers	Revenue commitment	No commercial viability	Utilities, hyperscalers
Funders	Capital provision	Project cannot be financed	Banks, equity, DOE LPO
Enablers	System integration, inputs	Project cannot operate	Grid operators, fuel suppliers
Delivery Actors	Physical execution	Project cannot be built	EPC firms, vendors
Legitimacy Actors	Social and political acceptance	Delays, opposition, escalation	Communities, NGOs
Strategic Sponsors	Policy and financial support	Reduced feasibility or scale	Governments

The key insight from this framework is that veto power is distributed and does not require formal authority. Project failure does not require a formal rejection from any institutional gatekeeper. It can result from the absence of any critical function. A project that satisfies every formal approval requirement can still fail if a delivery actor cannot execute, if a legitimacy actor succeeds in raising the cost of proceeding beyond what funders will accept, or if an enabler cannot provide required system integration on the project's timeline.

3.4 Classification by Decision Logic

The fourth classification groups stakeholders by the type of decision they make—the core objective, primary evaluation metric, and time horizon that govern their assessment of a project.

Table 4. Stakeholder Decision Logics

DECISION LOGIC	CORE OBJECTIVE	PRIMARY EVALUATION METRIC	TIME HORIZON	REPRESENTATIVE STAKEHOLDERS
Safety / Compliance	Prevent unacceptable harm	Regulatory standards, safety margins	Long-term	NRC
Commercial / Economic	Revenue viability	Cost, contract structure, offtake certainty	Medium—long	Utilities, buyers
Financial / Portfolio	Capital preservation and return	Risk-adjusted return, downside exposure	Medium	Banks, investors
System Reliability	Grid stability and adequacy	Integration feasibility, reliability metrics	Operational	ISOs, utilities
Political / Constituency	Maintain support, avoid backlash	Public response, coalition stability	Short—medium	Elected officials
Legitimacy / Social	Trust and fairness	Perceived impact, procedural legitimacy	Variable	Communities, NGOs
Strategic Positioning	Long-term advantage	Market position, optionality	Long-term	Corporates, governments

This classification is analytically the most important, because it aligns directly with how risk is defined. Each decision logic corresponds to a different category of risk and a different definition of acceptable evidence. These are not variations on a common framework. They are structurally different decision problems, applied to the same project.

3.5 Cross-Classification and Category Instability

Each framework captures part of the system while obscuring other parts. When applied together, they produce category instability: the same stakeholder occupies different positions depending on the analytical lens.

Table 5. Cross-Classification Example

STAKEHOLDER EXAMPLE	INSTITUTIONAL ROLE	SECTOR TYPE	FUNCTIONAL ROLE	DECISION LOGIC
Regulated Utility	Grid / Infrastructure	Quasi-public	Buyer	Commercial / Economic + System Reliability
Hyperscaler	Commercial (Buyers)	Private	Buyer / Strategic Sponsor	Reliability + Strategic Positioning
NRC	Regulatory	Public	Gatekeeper	Safety / Compliance
Bank	Financial	Private	Funder	Financial / Portfolio
Community	Public / Civic	Civic	Legitimacy Actor	Legitimacy / Social

Different stakeholders apply different frameworks simultaneously, which is why actors in one domain routinely misinterpret the behavior of actors in another. The developer who has satisfied regulatory requirements does not understand why financiers remain concerned. The financier who has modeled expected return does not understand why the community remains opposed. Each assessment is coherent within its own framework, but the frameworks are not translatable into one another. This is a structural failure of the system.

3.6 The Decision Domain as the Correct Unit of Analysis

Three conclusions follow from these classification systems together. First, no single framework is sufficient. Each captures part of the structure; none captures it whole. Second, stakeholder misalignment arises not only from conflicting objectives but from differences in how the system is conceptualized—actors are not fighting over the same thing, they are evaluating different things. Third, and most importantly: the appropriate unit of analysis is the decision domain, not the stakeholder. Stakeholders can change roles, enter and exit, be replaced by new institutional actors, or expand their engagement across domains. The domains themselves persist. A new buyer does not dissolve the commercial approval requirement, but rather, changes who must satisfy it. This distinction between actors and the domains they occupy is the foundation for the sections that follow.

Section 4 — Risk Fragmentation: Why "Risk" Is Not a Shared Variable

The most consequential misframing in large infrastructure development is the treatment of risk as a shared variable. Most proposed solutions—risk-sharing mechanisms, financial guarantees, insurance instruments, regulatory reform—rest on the assumption that risk can be identified at the system level, measured against a common standard, and reduced through targeted intervention. This assumption is wrong, which explains why so many such interventions fail to produce the intended result.

Risk in large infrastructure development is not a unified quantity. That would imply that safety-related risk is the same as financial risk. Instead, it is an institution-specific construct, defined by the objectives, constraints, and consequences that each stakeholder faces within its own domain. What appears as disagreement about risk magnitude is often a more fundamental mismatch: stakeholders are not evaluating the same thing at all.

4.1 Risk Domains Across Stakeholders

The same underlying event maps to different risk domains depending on which stakeholder is evaluating it. A construction delay is simultaneously a safety-schedule concern for the regulator, a cost-of-capital exposure for the lender, a counterparty performance question for the buyer, and a political liability for the sponsoring government. These are different variables, defined within different institutional frameworks.

Table 6. Stakeholder-Specific Risk Domains

RISK DOMAIN	CORE CONCERN	PRIMARY CONSEQUENCES	REPRESENTATIVE STAKEHOLDERS
Safety Risk	Probability and consequence of harm	Public health impact, regulatory failure	NRC, safety authorities
Regulatory Risk	Process uncertainty, delays, reinterpretation	Schedule extension, cost escalation	Developers, investors
Commercial Risk	Revenue uncertainty, performance variability	Contract failure, reduced returns	Utilities, buyers
Financial Risk	Capital loss, default probability	Loss of investment, portfolio impact	Banks, equity investors

RISK DOMAIN	CORE CONCERN	PRIMARY CONSEQUENCES	REPRESENTATIVE STAKEHOLDERS
System Reliability Risk	Grid stability and adequacy	Service disruption, system instability	ISOs, utilities
Political Risk	Public and electoral response	Policy reversal, loss of support	Policymakers
Social / Legitimacy Risk	Trust, fairness, perceived harm	Opposition, litigation, delays	Communities, NGOs
Strategic Risk	Sector exposure, long-term positioning	Entry/exit decisions, capital allocation shifts	Corporates, financial institutions

Risk fragmentation is reinforced by differences in measurement and evidence. Regulatory institutions evaluate risk through formalized methodologies requiring probabilistic analysis and documentation. Financial institutions use models of expected loss and portfolio-level exposure. Commercial actors focus on contract structure and counterparty quality. Political and social actors rely on perception, precedent, and anticipated public response. Evidence that resolves risk in one domain may have no relevance in another. Demonstrating compliance with safety standards does not address financing risk. A strong power purchase agreement does not resolve regulatory uncertainty. A favorable financial model does not address public trust.

Risk cannot be resolved once and transferred across stakeholders. It must be independently resolved within each domain, often using different forms of evidence and different standards of sufficiency.

4.2 Risk Translation Failure

The same word—"risk"—refers to different constructs across domains. This produces systematic communication failures. One stakeholder's demonstration of acceptable risk does not resolve another's concern, because the underlying concepts are different.

Typical patterns:

- A developer demonstrates regulatory compliance and characterizes the project as low risk; investors remain concerned about cost and schedule exposure that safety compliance says nothing about.
- A financial institution declines participation not due to expected project loss but due to strategic or reputational considerations that are not captured in any project-specific risk model.

- A community rejects a project despite technical safety assurances, because the community's concern is not modeled risk but perceived fairness and institutional trust.

In each case, the issue is a mismatch in what is being evaluated rather than a disagreement over a shared metric. These represent structural differences in how each stakeholder defines acceptable outcomes that cannot be resolved through "better" communication or more data.

4.3 Risk Interaction and Compounding

Risk domains do not operate in isolation. Uncertainty in one domain propagates into others, often amplifying overall exposure in ways that are invisible from within any single domain:

- Regulatory delay extends capital exposure and raises interest costs, converting timeline uncertainty into financial risk.
- Supply chain uncertainty increases both commercial and financial risk through schedule variability, converting delivery uncertainty into revenue risk.
- Political opposition extends timelines, creating cascading exposure across regulatory, financial, and commercial domains simultaneously.
- Macroeconomic shifts increase financial risk through capital costs and commercial risk through competitiveness, often at the same time and for the same reason.

These interactions produce compounding effects—risks that reinforce one another rather than remaining bounded within their domain of origin. A project can move from acceptable to unacceptable across multiple domains in sequence from a single initial disturbance, and this propagation is not visible to any single stakeholder managing risk within their own domain.

4.4 Risk Allocation and the Transformation of Risk Types

Because risk is defined independently across domains, it is actively allocated across stakeholders through policy design, contract structure, and institutional arrangement. It is rarely eliminated. It is redistributed—and in being redistributed, it is transformed in type. And risk that is manageable in one domain may become destabilizing in another. Several patterns are recurrent.

Policy-driven risk transfer. Government loan guarantees and cost recovery mechanisms reduce financial risk for developers and investors by transferring exposure to the public sector. Financial risk becomes political risk: cost overruns or project failure that would previously have been a private loss become a public liability with electoral consequences, affecting the durability of the policy support that made the project viable in the first place.

Contractual risk concentration. EPC contract structures often allocate contingency and performance risk to a single delivery actor to provide cost certainty to developers and financiers.

This concentration creates counterparty risk: if the contractor lacks the balance sheet to absorb losses, financial distress or failure transfers the risk back to the project—and to the other stakeholders who had assumed it was contained. This dynamic played out at Vogtle and V.C. Summer, where Westinghouse's inability to absorb fixed-price losses ultimately produced contractor bankruptcy, project disruption, and renewed exposure for owners and financiers.

Regulatory requirements and supply chain risk. Nuclear-specific quality assurance requirements shift risk to suppliers by increasing qualification cost, liability exposure, and documentation burden. This reduces regulatory risk for operators while discouraging supplier participation. The result is a supply chain constraint that emerges not from external market conditions but from a risk allocation decision embedded in the regulatory structure.

Demand-driven risk reallocation. The entry of hyperscalers as direct buyers reduces revenue uncertainty for developers and improves commercial viability. However, the compressed timelines they require increase exposure to infrastructure and supply chain constraints. Commercial risk is reduced; execution risk and system integration risk increase. The risk is not eliminated—it is shifted toward the domains least capable of absorbing acceleration.

This implies a general principle: risk mitigation in one domain is frequently risk transfer to another domain, often with transformation in type. Effective risk allocation requires matching risk types to stakeholders with the structural capacity to absorb them.
Misallocation—concentrating risk in actors without sufficient capacity, or transferring risk in ways that convert manageable exposure into systemic vulnerability—increases the probability of system-level failure.

4.5 Supply Chain Risk as a Hybrid Constraint

Supply chain risk illustrates the interaction between institution-specific risk and structural constraint in a way that generalizes beyond nuclear.

Nuclear supply chain limitations—constrained manufacturing capacity, limited qualified suppliers, long lead times—are partially produced by the regulatory and institutional structure. Qualification requirements increase supplier entry cost. Liability and compliance frameworks create risks that differ from other industrial sectors. Limited and uncertain demand discourages investment in qualified capacity. This produces a coordination problem: suppliers require credible demand to justify investment, while projects require supply capacity to proceed. Neither side can fully resolve the other's uncertainty first.

The same dynamic operates in sectors with none of nuclear's specific regulatory barriers. Gas turbine manufacturers are currently expanding production capacity in response to strong demand from data center development, industrial reshoring, and grid replacement. They remain sold out past 2030, even accounting for expansion that is actively underway. The constraint, in this case, is rational risk aversion: manufacturers will not build capacity beyond what committed orders justify, because demand that cannot be credibly sustained over the investment horizon does not warrant the capital risk of excess capacity. Projects that need equipment cannot wait for the manufacturers to become confident in demand. The co-evolution of supply constraints and demand uncertainty produces a binding limit that does not originate with any single stakeholder and cannot be resolved by any single domain.

This is the correct structure of a quasi-exogenous constraint—produced by the interaction of stakeholder decisions and institutional design, experienced by individual actors as an external limit, and addressable only through coordinated changes across multiple domains.

4.6 Hyperscalers and the Entry of a New Risk Class

The emergence of hyperscale technology firms as direct energy market participants introduces a risk profile that differs from any prior buyer class.

Traditional utilities prioritize cost stability, regulatory approval, and resource adequacy. Their risk tolerance reflects the constraints of regulated cost-of-service frameworks and the political accountability of rate-setting processes. Hyperscalers operate under different constraints. Their direct operational dependence on continuous power means that reliability risk and delay risk—not cost risk—govern their decision calculus. Deferred capacity directly affects revenue and competitive position. Cost uncertainty is acceptable; schedule uncertainty is not.

Hyperscalers therefore accept risks that utilities have historically avoided, particularly construction risk and schedule risk on new projects. Their entry into energy markets does not reduce system-level risk, but instead redistributes and reweights it. By accelerating demand formation and compressing commercial timelines, their participation brings forward bottlenecks in interconnection, transmission, supply chain, and regulatory throughput.

Section 5 — Non-Fungible Approvals: Why Solving One Constraint Does Not Solve the System

Large infrastructure projects are frequently described as a sequence of approvals—licensing, then financing, then construction—that must be obtained in order to advance the project toward completion. This description incorrectly implies that progress is accumulative: each approval achieved represents progress toward realization.

Project realization is not the accumulation of sequential approvals. It is the simultaneous satisfaction of independent conditions—conditions that do not substitute for one another, cannot be partially offset against one another, and do not collectively advance the system unless all hold at once. A project that has satisfied every condition but one is not almost finished. It is at zero, because the unsatisfied condition is necessary.

These approvals are non-fungible. The term is precise: non-fungible means not interchangeable, not substitutable, not transferable in value across domains. Regulatory approval cannot offset commercial failure. Financial commitment cannot substitute for social acceptance. Technical clearance cannot remedy political obstruction. Each domain resolves a distinct question, and no other domain can answer it.

Table 7. Non-Fungible Approval Types and Associated Risk Domains

APPROVAL TYPE	CORE QUESTION	DECISION AUTHORITY	PRIMARY RISK DOMAIN	FAILURE MODE
Legal / Regulatory	Does the project meet statutory requirements?	Regulators, permitting agencies	Safety / Regulatory	Cannot proceed legally
Commercial	Is there a viable revenue pathway?	Buyers	Commercial	No economic basis
Financial	Is the project investable?	Capital providers	Financial	Cannot secure capital
Technical / Infrastructure	Can the project operate within the system?	Grid operators	System Reliability	Cannot connect or operate

APPROVAL TYPE	CORE QUESTION	DECISION AUTHORITY	PRIMARY RISK DOMAIN	FAILURE MODE
Political	Is the project supported or tolerated?	Policymakers	Political	Policy obstruction
Social / Legitimacy	Is the project acceptable locally?	Communities	Social / Legitimacy	Opposition and delay

This structure follows directly from risk fragmentation. Risk is defined independently across domains, so the decisions that resolve those risks cannot be aggregated into a single approval. For example, regulatory approval resolves safety and compliance risk. Financial approval resolves capital exposure and return risk. These are not the same risk. They cannot be resolved by the same decision. Joint satisfiability—the requirement that all conditions hold simultaneously—is therefore not an organizational design choice. It is a structural consequence of risk fragmentation across independent institutional domains.

Parallel evaluation dynamics and interactions across domains compound the problem. Approvals are evaluated simultaneously by different stakeholders under conditions of uncertainty. Financial institutions assess projects based on anticipated regulatory outcomes. Buyers evaluate commitments based on projected cost and delivery timelines. Developers advance licensing without certainty about financing or demand. Each stakeholder makes conditional decisions based on expectations about other domains. These interactions produce a system of contingent assessments without centralized coordination—everyone is waiting for everyone else, conditioned on assumptions about what everyone else will decide.

One conjunctive condition across independent approvals

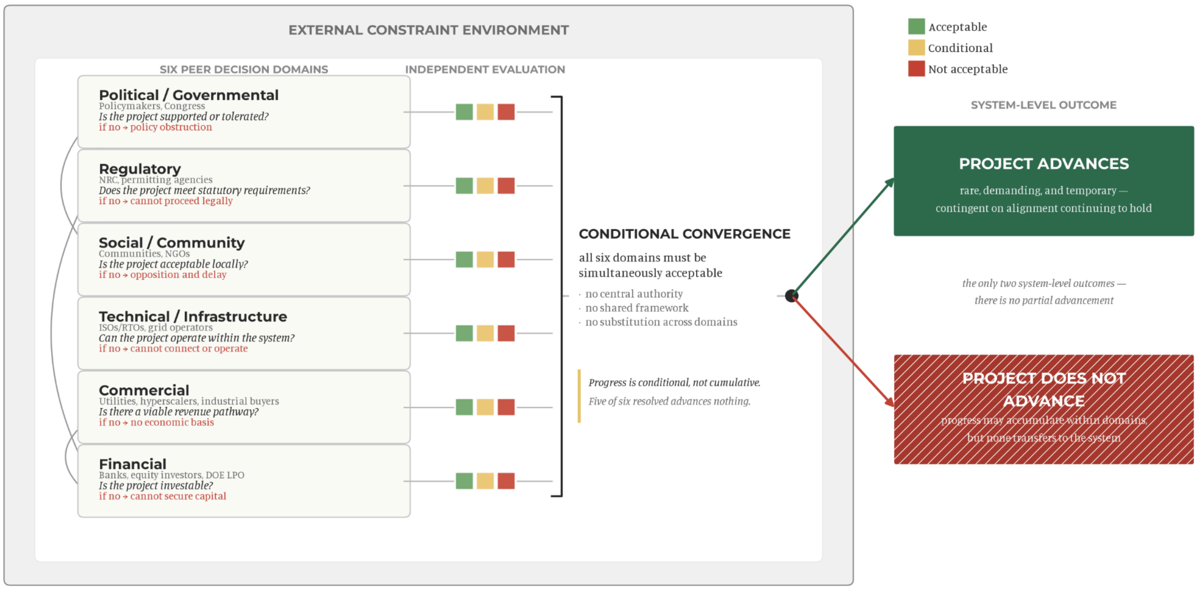

Figure 1. *Distributed Approval and Constraint System for Nuclear Project Realization. Interactions across layers are shown conceptually. The domains shown are representative, not a fixed set.*

Project progress is conditional rather than cumulative, as illustrated in Figure 1. Advancement in one domain produces no system-level progress unless other domains converge simultaneously. External constraints—macroeconomic conditions, supply chain capacity, infrastructure limits—act across all domains and can prevent convergence even when individual approvals are attainable.

The interaction between non-fungibility and risk allocation is direct. Efforts to secure approval in one domain typically redistribute risk to others, often in ways that change what is required in those domains. Government financial support reduces financing risk while increasing political exposure. Contract structures that concentrate execution risk in a supplier reduce developer and funder exposure while creating counterparty risk that can destabilize the whole project if the supplier cannot absorb what they have been allocated. These reallocations determine where risk resides and whether it is held by actors with the capacity to absorb it.

The characteristic outcome of this structure is partial progress without realization: projects that advance substantially within one or more domains—completing licensing, securing preliminary financing, establishing commercial agreements—and then fail because a different domain never converged. This is the most common failure mode in large nuclear project development. It is frequently interpreted as a coordination failure. Within this framework, it is expected. It is what a system of non-fungible approvals evaluated in parallel under uncertainty produces.

Project realization, then, depends on achieving simultaneous acceptability across independent decision domains. Even when that condition is satisfied, realization remains contingent on whether prevailing constraints permit execution.

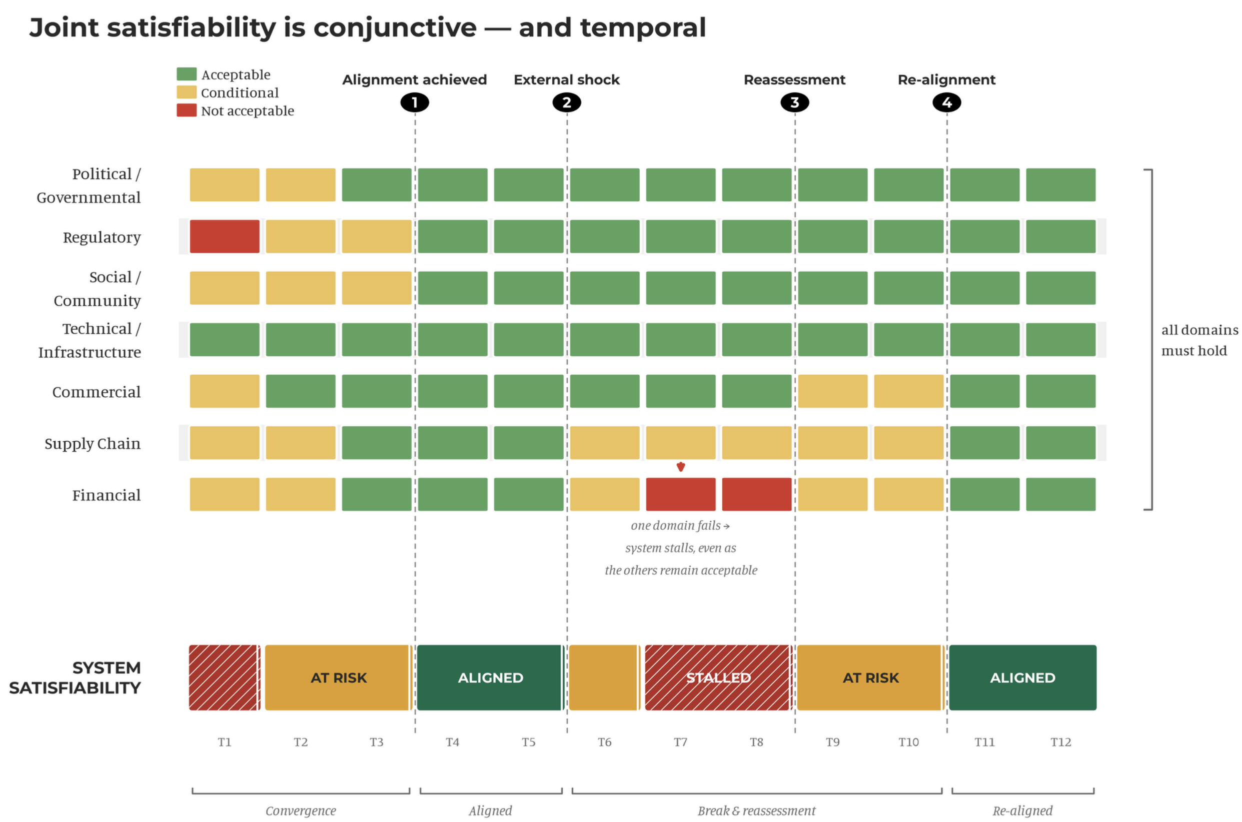

Figure 2. *Illustration of joint satisfiability as temporal equilibrium.*

The temporal dimension of joint satisfiability is illustrated in Figure 2. Each domain must reach an acceptable threshold simultaneously—convergence in one domain while others remain outside their thresholds produces no progress. The system feasibility track at the bottom shows that even brief failure in a single domain resets the conjunctive condition.

Section 6 — External and System-Generated Constraints: Feasibility Beyond Approval

The preceding sections establish that large infrastructure project realization depends on distributed decision-making, fragmented risk definitions, and the simultaneous satisfaction of non-fungible approval conditions. While these elements describe how projects are evaluated and authorized, they do not fully determine whether they can be executed.

Approval is a necessary condition for project realization, but it is not a sufficient one. Execution depends on feasibility—whether the project can be physically, economically, and institutionally delivered under prevailing conditions. Feasibility is shaped by constraints that operate alongside the approval system: some originate outside the control of project stakeholders, while others are generated by the interaction of institutional structures, incentives, and prior decisions. Any of these can prevent execution even when all approval conditions have been satisfied. Section 6.1 classifies these constraints by origin, while Section 6.2 explores how they are generated, developing in depth the mechanism of risk allocation.

The distinction matters because approval and feasibility are governed by different systems that do not automatically align. A project fully approved under one set of macroeconomic conditions may become unfinanceable when those conditions shift. Similarly, a project authorized assuming supply chain availability may become undeliverable as capacity tightens. Approval establishes authorization; feasibility determines whether that authorization can be converted into a built project.

6.1 Constraint Typology: External and Quasi-Exogenous

Constraints differ in their origin—specifically, where they come from and whether the stakeholder system produced them. Three categories describe the standing conditions under which a project operates; distinguishing them is essential because they each respond to different interventions.

Table 8. Constraint Categories Affecting Project Feasibility

CONDITION EXAMPLE	CONSTRAINT TYPOLOGY	MECHANISM	PRIMARY IMPACT	EXAMPLE EFFECTS
Macroeconomic	External	Cost of capital, inflation, credit conditions	Financial viability	Projects become unfinanceable under rising interest rates

CONDITION EXAMPLE	CONSTRAINT TYPOLOGY	MECHANISM	PRIMARY IMPACT	EXAMPLE EFFECTS
Supply Chain	Quasi-Exogenous	Capacity limits, qualification requirements, lead times	Execution feasibility	Delays from limited qualified suppliers; sold-out capacity
Infrastructure	Hybrid	Physical system limits, interconnection queue, transmission	Integration feasibility	Grid integration delays; transmission unavailability

External constraints are imported into the system rather than being created by stakeholders within it. These are conditions to which developers, customers, suppliers, regulators, and investors must respond. Macroeconomic conditions are the clearest example: interest rate environments, inflation, and credit availability are determined by forces that no project stakeholder controls, yet they can shift substantially over the decade-long timelines of large infrastructure projects. A project financed under one rate environment may become unviable under another—not because of any internal change, but because the operating conditions no longer match those of the approval stage.

Quasi-exogenous constraints are distinct. While they appear external to any individual actor, they are generated by the deployment system as a whole. No single developer creates a shortage of qualified suppliers, nor does an individual customer create component scarcity; rather, these constraints emerge from the aggregate behavior of stakeholders operating under uncertainty. They are not simply imposed from the outside; they are produced by the system's own incentives, sequencing, and prior decisions.

Hybrid constraints straddle the boundary between external and quasi-exogenous. They may begin as conditions beyond stakeholder control but become more binding due to choices made within the deployment system. Grid interconnection queues, for instance, reflect the aggregate requests of developers responding to demand signals that grid operators did not create. Similarly, transmission constraints reflect investment decisions based on prior forecasts that current demand has rendered obsolete. These are external to any individual project yet are produced by the same institutional system in which projects participate—shaped by both external shocks and system-generated dynamics.

6.2 Constraints as Products of Risk Allocation

The typology in Section 6.1 classifies constraint categories by origin. The production of system-generated (quasi-exogenous) constraints, however, is a question of mechanism. While a range of such mechanisms exists, this report focuses on the most prominent ones rather than an exhaustive catalog.

The public serves as one example, creating constraints along two different paths. Most often, the public lacks a direct lever matched to its objectives, so its influence projects indirectly: negative sentiment hardens political opposition, makes buyers reluctant to attach themselves to a contested project, and leads lenders to price a higher probability of delay or cancellation. In each case the public is moving the stakeholders who hold levers, not acting on the project itself. Where the public does hold a lever of its own, such as environmental litigation, it acts on the project directly, and often uses that lever as a proxy. The underlying objection is frequently to the project's existence rather than its environmental impact, and the litigation is the one procedural handle available: its purpose is often to delay or exhaust the project. By either path, a group with little formal standing in the decision system produces a binding constraint.

Risk allocation is another such mechanism, and the one this section develops, because it is most consistently misunderstood and improperly applied. Conventional treatments often speak of "risk" as a single quantity to be reduced or retired without asking where it is actually allocated. But risk is not fungible. As argued in Section 4, what counts as risk differs across domains; therefore, the question governing outcomes is not merely the magnitude of risk, but where it is placed and what it becomes upon arrival. Allocation, not just magnitude, is a critical variable.

Risk allocation is the reassignment of a risk from one actor to another—by transferring it, as a fixed-price contract transfers cost risk to a contractor, or by internalizing it, as a build-own-operate developer does by taking construction, financing, and operation onto a single balance sheet. Two properties make its consequences difficult to anticipate.

The first follows from non-fungibility. Allocating risk across layers does not move a fixed quantity between ledgers; it converts the risk into a different category that the receiving layer defines and weighs on its own terms. Cost-overrun risk transferred to a contractor becomes counterparty risk once the contractor's balance sheet is at stake. Financing risk transferred to the public through a guarantee becomes political risk once a cost overrun is a public liability with electoral consequences. Regulatory-compliance risk transferred to suppliers through a qualification mandate becomes supply-chain risk once qualified vendors grow scarce. The risk does not shrink in transit; it changes form, and it often lands in a domain less able to absorb it.

The second property is that a transfer requires an allocator but not necessarily a consenting recipient. Some transfers are negotiated and accepted: a contractor signs a fixed-price contract, a lender prices a premium into its rate. Other transfers are imposed on parties with no seat at the table—a government shifting cost-overrun risk onto taxpayers, a utility passing it to ratepayers, a party using environmental litigation to impose delay risk on a developer it has no commercial relationship with. Consent or not, the risk becomes the recipient's, and the recipient's response—whether it be withdrawal, opposition, a higher cost of capital, a lawsuit—becomes a new constraint on the project.

Together these properties explain why an allocation made in one place surfaces as a binding constraint in another, along a consistent sequence. A risk exists in one domain. A stakeholder reassigns it rather than resolving it. In moving, it changes type and lands on another party in another domain, sometimes consenting, sometimes not. There, transformed and frequently unpriced, it becomes binding, often in a domain that was not binding before. By the time it binds it resembles nothing about the decision that produced it: different form, different owner, later date. The counterparty constraint that appears when a fixed-price contractor cannot absorb its losses does not announce itself as the downstream consequence of a cost-certainty decision made years earlier—but that is what it is.

Table 9. Risk Transfer and Constraint Formation

CASE PATTERN	INITIAL STRATEGY	TRANSFER MECHANISM	RESULTING CONSTRAINT	SYSTEM EFFECT
Government financial support	Reduce financing risk	Risk shifted to public sector	Political constraint	Reduced policy stability
EPC risk concentration	Reduce developer exposure	Risk concentrated in contractor	Counterparty constraint	Project delay or failure
Nuclear QA requirements	Reduce regulatory risk	Risk shifted to suppliers	Supply chain constraint	Reduced vendor capacity
Hyperscaler demand	Reduce commercial risk	Timeline pressure on system	Infrastructure constraint	Bottlenecks in delivery
Fixed-price contracts	Reduce cost uncertainty	Risk shifted to delivery actors	Execution constraint	Elevated failure probability

Each row of Table 9 traces one such sequence. Government financial support reduces financing risk by converting it into political risk: a cost overrun that would have been a private loss becomes a public liability, and the constraint reappears as reduced durability of the policy support the project depends on. Fixed-price EPC structures reduce developer and financier exposure by concentrating execution risk in the contractor; if the contractor cannot absorb it, the constraint reappears as counterparty failure that returns the risk to the project in a more damaging form (e.g., supplier bankruptcy). Nuclear quality-assurance requirements reduce regulatory risk by shifting qualification cost and liability onto suppliers; the constraint reappears as reduced vendor capacity, which developers who never made the trade encounter as part of the thin supplier base described in Section 6.1—an internal allocation that has hardened into a structural condition others meet as quasi-exogenous.

The practical implication is that because an allocated constraint has a source, it can be addressed at its source, but only if the analysis traces it there. A constraint treated only where it surfaces, without tracing the allocation decision that produced it, cannot be resolved, only relocated: the system will reconstitute it in whichever domain next receives the transferred risk. Section 8 develops case illustrations of exactly this failure, including interventions that addressed a visible constraint while leaving untouched the allocation decision that generated it.

6.3 Dynamic Feasibility

Feasibility is a dynamic condition that evolves as macroeconomic, supply chain, and infrastructure environments change—sometimes during execution and after major commitments have been made.

Projects are evaluated under one set of conditions and must be executed under conditions that may differ materially. A financial structure built on assumptions about interest rates, supply chain availability, and grid capacity may encounter all three assumptions invalidated before the project reaches commercial operation. This temporal dimension parallels the alignment dynamics described in Section 7: just as stakeholder alignment can degrade asymmetrically as conditions change in individual domains, feasibility can degrade asymmetrically as constraints tighten in domains that were not binding at the time of authorization.

The interaction of approvals and constraints produces a characteristic pattern that explains many of the cases in the large-project performance literature: projects that achieve full stakeholder alignment and then fail when constraints become binding during execution. This outcome is not a simple coordination failure that better planning could have prevented. It reflects the fundamental structure of a system in which authorization and execution depend on different conditions, governed by different actors, under different timelines.

Section 7 — Alignment and Conflict: Structural Limits to Coordination

The preceding sections establish that decision authority is distributed, risk is defined differently across stakeholder domains, and feasibility is bounded by constraints that can shift independently of approval outcomes. Taken together, these conditions imply that achieving project advancement is structurally constrained—and that the constraint is architectural, not behavioral.

The prevailing framing of coordination problems in large infrastructure treats misalignment as a product of disagreement that better communication, institutional design, or incentive alignment could resolve. This framing is misleading. It treats a structural condition as an informational problem and prescribes solutions—stakeholder engagement, consensus processes, benefit-sharing mechanisms—that address the symptom while leaving the structure intact.

In practice, alignment requires something more limited and more demanding: joint acceptability across independent decision systems. Each stakeholder must independently determine that the project meets its own criteria for action, under its own constraints, using its own definition of acceptable risk. A regulator may approve a project on safety grounds while a capital provider evaluates it on risk-adjusted return and a buyer assesses it on reliability and schedule certainty. These decisions are not coordinated through a common framework. They are conditionally aligned outcomes across structurally different evaluation systems. When they align, it is because each stakeholder independently reached an acceptable threshold in their own domain at the same time.

This produces a conjunctive requirement: failure in any single domain prevents realization regardless of success in all others. The system therefore does not benefit from partial alignment. Every domain but one at an acceptable threshold advances nothing if the remaining one is outside it.

7.1 Why Alignment Is Structurally Difficult

The difficulty of achieving joint acceptability arises not primarily from disagreement but from three structural conditions.

First, **risk translation failure.** Risk is defined differently across domains (Section 4), and evidence that resolves risk in one domain does not translate to another. A project that is low risk to a regulator may remain unacceptable to a financier—not because one lacks information, but because both are evaluating a different type of risk entirely. The domains must each be resolved on their own terms.

Second, **redistribution fragility**. Risk allocation decisions made to achieve alignment in one domain often create new vulnerabilities in others. As Section 4.4 shows, alignment achieved through risk redistribution is often internally fragile: it holds only as long as the receiving stakeholder can absorb what has been allocated to them. When that capacity is exhausted, the risk returns to the system in transformed form, often in a domain that was not part of the original calculation.

Third, **temporal instability**. Large projects unfold over time horizons during which macroeconomic conditions, supply chain capacity, political environments, and infrastructure availability all evolve. As established in Section 6, these changes operate independently of stakeholder authorization decisions and can alter feasibility after alignment has been achieved. Alignment is a temporal equilibrium that must persist across multiple independent domains under conditions that continue to change.

Each stakeholder updates its assessment based on conditions in its own domain, without centralized oversight of the system as a whole. Alignment can therefore degrade asymmetrically: a financing condition changes and capital providers withdraw, even if regulatory and commercial conditions remain intact. An infrastructure constraint tightens and commercial viability erodes, even if demand remains strong. Projects can move forward for years and then collapse when a single domain becomes binding. This is the expected behavior of a system in which alignment must be maintained simultaneously across independent domains under evolving conditions.

7.2 Demand Acceleration and System Constraints

Recent electricity demand growth illustrates the limits of single-domain acceleration. Strong commercial demand signals—from data centers, industrial reshoring, and grid replacement needs—have improved commercial and, in many cases, financial alignment across multiple generation technologies. This is meaningful. It is also insufficient.

Supply chains have not responded at the pace that commercial demand would suggest. Gas turbine manufacturers will not build capacity beyond committed orders, because demand that cannot be credibly sustained does not warrant the capital risk of excess capacity. This is the quasi-exogenous constraint structure from Section 6.1, operating in one of the most mature, least-regulated industrial supply chains available. The co-evolution of supply constraints and demand uncertainty is a structural feature of capital-intensive supply systems operating under demand uncertainty, and is not limited to just nuclear, or gas turbine manufacturing.

In parallel, regulatory timelines remain set by their own processes, infrastructure constraints are governed by interconnection queues and transmission investment cycles that predate the current demand environment, and political conditions are shaped by electoral calendars and institutional dynamics that accelerate on their own terms. Commercial alignment has improved. The other domains have not correspondingly accelerated. Faster convergence in one domain does not produce faster convergence in others. It exposes the gap between where commercial alignment has arrived and where the other domains remain.

7.3 Structural Conflict

Stakeholder conflict in large infrastructure is typically analyzed as a problem of interests—actors who want different outcomes, whose opposition can be resolved through negotiation, benefit-sharing, or institutional design. This framing is accurate in some cases and seriously misleading in others.

Where conflict arises from genuinely incompatible definitions of acceptable risk, resolution is not available through negotiation. A community whose concern is procedural fairness and institutional trust is not evaluating the same thing as a regulator whose concern is modeled radiological risk. A financial institution whose concern is portfolio-level exposure to a new-build category is not evaluating the same thing as a developer whose concern is project-specific construction feasibility. These are different problems that do not resolve through a common answer.

Treating structural conflict as a coordination failure—as a problem of information, communication, or incentive design—misdiagnoses the condition and leads to interventions that address the symptom. Stakeholders operating under different constraints, evaluating different risks, and bearing different consequences will reach different conclusions about the same project. This is the expected behavior of independent decision systems applied to a shared subject.

7.4 The Falsifiability Condition

The framework developed in this report makes a structural claim: in systems with genuinely distributed decision authority, alignment is inherently transient, coordination is inherently fragile, and project realization is structurally uncertain even when alignment appears to have been achieved. This claim is falsifiable.

It would be refuted by a clear case in which decision authority is genuinely distributed across independent stakeholders and yet alignment is durable and coordination is not fragile. No accumulation of confirming cases establishes the claim—structural regularities cannot be proven by enumeration—but a single well-documented counterexample would disprove it.

The framework predicts a specific attenuation: the condition weakens where authority consolidates. A vertically integrated utility that internalizes generation, financing, and delivery faces fewer independent decision domains. A single-buyer or state-directed program with few independent approval requirements exhibits less transient alignment because fewer independent decisions must converge simultaneously. These are predicted attenuations, not demonstrated ones. This report does not supply the disconfirming case.

The claim is strongest where decision authority is most distributed, and it has intensified as the energy sector has de-concentrated across independent regulators, financiers, developers, suppliers, and buyers.

Section 8 — System Implications: Why Interventions Fail at the System Level

The system described in this report is defined by distributed authority, non-fungible approvals, fragmented risk, and binding constraints. These characteristics determine outcomes jointly. Large infrastructure project development is not an optimization problem. It is a coordination problem under constraint. This distinction has a direct implication: single-domain interventions, however well-designed within their domain, cannot resolve the system. They can only alter the location and timing of the binding constraint.

This explains a persistent pattern in large infrastructure policy: interventions that are analytically coherent and technically sound within their domain fail to produce system-level outcomes. The analysis that produced them is incomplete. It correctly identifies a constraint and correctly characterizes a solution for that constraint. It fails to trace where the constraint will re-emerge when displaced. The displaced constraint appears in a different domain, at a later stage, attributed to a different cause. The original intervention is assessed as successful. The system has not changed.

Risk allocation is the central design variable precisely because it is through risk allocation that constraints are produced and displaced. Risk transferred to an actor without the capacity to absorb it re-emerges as instability—through contractor failure, policy reversal, supply chain contraction, or financing withdrawal. Effective system design requires not only identifying where risk currently resides, but mapping how it moves when displaced and whether the receiving stakeholder can absorb it without reconstituting the constraint elsewhere.

8.1 Solving the Problem at the Wrong Level

Cost overruns at the Vogtle project substantially increased perceived risk among state public utility commissions (PUCs), which are responsible for protecting ratepayers from imprudent expenditure. In response, proposed federal policy interventions have focused on partially offsetting potential cost overruns on future projects, with the objective of increasing PUC willingness to authorize new nuclear development.

This approach correctly identifies the PUC as a constraint on the political and commercial approval process for new projects. The intervention addresses the wrong problem.

Vogtle's overruns were produced by construction complexity, supply chain limitations, first-of-a-kind manufacturing challenges, and an EPC contract structure that concentrated risk in a contractor—Westinghouse—that lacked the financial capacity to absorb it. When Westinghouse declared bankruptcy in 2017, execution risk returned to the project owners and ultimately to ratepayers in a far more damaging form than any overrun-sharing mechanism could have anticipated. The intervention addresses the symptom (PUC risk aversion) rather than its source (execution risk in nuclear construction, which is a delivery and supply chain problem, not a regulatory problem).

Evidence of the misdiagnosis appears in the PUC responses themselves. Regulators in several states have indicated that partial overrun protection is insufficient to materially change their position, with some suggesting that only full cost protection would do so. This is not an unreasonable position given the evidence. But full federal cost protection would shift risk from ratepayers to federal taxpayers without improving the delivery capability that produced the overruns in the first place. The political domain would receive an intervention. The supply chain and delivery domains would remain unchanged. V.C. Summer confirms the pattern: cost escalation and execution risk produced project abandonment despite full regulatory approval and substantial financial commitment. The binding constraint was upstream of the domains being addressed.

8.2 Adding Decision Points

A distinct failure mode occurs when an intervention does not merely shift risk across existing domains but creates a new procedural surface through which the project must be continuously managed. This is emergent procedural stratification: a new decision layer arises from the process itself, independent of the project's original approval architecture.

The V.C. Summer environmental review illustrates the mechanism. The draft environmental impact statement found that noise from construction and operation would be intermittent and that associated impacts would not be significant.[2] EPA's review did not contest that finding. Nevertheless, EPA recommended that a community advisory group be established for local residents, with periodic meetings during project development and operation.[3]

The proportionality is instructive. The review found that noise at sensitive receptors approximately one mile from the site would attenuate below 50 dBA—approximately the ambient level of a quiet office. The recommended institutional response to a non-significant impact at that level is a continuing governance process with no defined scope, no defined duration, and no defined threshold for termination. A technically limited environmental finding generated an open-ended relational obligation.

The structural consequence is durable. Once established, the advisory process converts a discrete technical question—is the noise impact significant?—into a continuous governance relationship: is the project adequately engaging, reporting, responding, and adapting over time? The original question has already been answered. The new process generates new questions indefinitely.

The V.C. Summer recommendation is one instance, but the mechanism it illustrates is not incidental to that project. Wherever a non-significant technical finding can justify a continuing engagement process, the same procedural layer can emerge. Community benefits agreements (CBAs) formalize this mechanism deliberately: negotiated commitments through which local assent, non-opposition, or benefit delivery becomes part of the project's operating conditions.[4—6] Their use expanded significantly when the Department of Energy began requiring Community Benefits Plans for nearly all funding opportunities and loan applications.

Where V.C. Summer shows emergent procedural stratification, CBAs create explicit contractual structures through which local acceptance becomes conditional and negotiated. Once formalized, local assent is no longer a binary condition. It is a variable—subject to renegotiation as conditions change, as expectations evolve, and as stakeholders recognize that their continued cooperation has retained value beyond the initial commitment. The decision shifts from whether a community supports a project to what terms it will accept and for how long. This pattern is visible across energy sectors: in California's Diablo Canyon coastal process, in offshore wind development, and in transmission siting proceedings where broad regional support has not prevented extended local negotiations over compensation and governance.[4, 5, 7, 8, 9]

Formalizing assent introduces an open-ended process into the project structure that appears external to other stakeholders. Financial actors must evaluate not only whether community agreements exist but the uncertainty associated with their negotiation, scope, evolution, and enforceability—uncertainty that does not have a clear stopping rule and therefore does not parameterize easily into standard financial risk models. The mechanism intended to reduce social or political risk creates a new category of process risk.

8.3 Using the Available Stratum

The preceding cases involve actors operating within domains where they already hold recognized authority. A structurally distinct pattern arises when an actor holds no authority in the domains that govern a project's merits—financing, technical review, commercial offtake—but does hold standing in one adjacent domain: judicial review of agency procedure.

Environmental review under the National Environmental Policy Act creates such a domain. Because NEPA governs the adequacy of an agency's procedural record rather than the substantive wisdom of a project decision, it furnishes a decision point that actors outside the project's core approval domains can reach and contest. The statute is procedural. Its availability as a litigation vehicle is therefore not conditional on having a substantive role in the project's authorization.

An analysis of more than 1,400 federal NEPA rulings between 2013 and 2022 found that energy projects spent a median of three years between final agency approval and final judicial decision, regardless of outcome.[10] The delay is the operative effect: a completed approval is converted back into an open question for the duration of the proceeding. The litigation stratum is used most intensively where projects are largest—full environmental impact statements, representing approximately one percent of all NEPA reviews, accounted for 37 percent of district court and 42 percent of circuit court energy project rulings. The forum is reached most often precisely where the most decision domains must converge.

The distribution of litigants confirms the structural character of this behavior. Environmental organizations appeared in roughly three-quarters of judgments; a single organization appeared in approximately 28 percent of energy project rulings; three organizations accounted for nearly half of wildfire-risk decisions. A small set of actors repeatedly using a single stratum is exactly the behavior the framework predicts: when an actor's effective authority is concentrated in one forum, that forum becomes its instrument, regardless of the domain in which its underlying concern originates.

The effect does not depend on the challenger prevailing. Courts found a genuine agency flaw in approximately 26 percent of rulings. While any case proceeds, however, its uncertainty propagates across the other domains: developers defer expenditure, investors discount the project or wait for resolution, agencies prepare supplemental analysis, and buyers hesitate to commit. A procedural challenge in the judicial stratum becomes a financial, commercial, and political constraint in all the others. The interval of uncertainty is what the other stakeholders must price, not the eventual judgment.

This reframes a debate usually conducted in terms of good and bad faith. Whether a litigant is motivated by genuine environmental concern or is using NEPA instrumentally to delay a project it opposes for other reasons, the structural effect is identical: an actor without authority in the governing domains acquires, through the one domain it can reach, the ability to suspend decisions made in domains it does not control. The behavior is rational given the architecture. A non-substitutable decision surface operates whether or not any given challenge is meritorious.

The underlying dataset covers only cases producing written judicial opinions, understating settlements and dismissals; the concentration it shows is therefore a conservative measure.

8.4 Correct but Incomplete Problem Framing

The trajectory of hyperscaler participation in energy markets illustrates a fourth pattern: an intervention correctly identifying a real constraint but incompletely defining the problem, discovering the system's full structure through successive encounters with the domains it had not yet entered.

Hyperscalers entered the energy market with a specific diagnosis: utility procurement was too slow, too conservative, and insufficiently responsive to rapid load growth. Traditional utility planning processes—governed by long integrated resource planning cycles, conservative demand forecasts, and regulatory approval timelines—could not accommodate the scale or speed of data center power demand. The proposed solution was direct procurement: long-term bilateral power purchase agreements with developers, bypassing the utility planning process to create a parallel commercial pathway for new generation.[11,12]

The diagnosis was accurate within its domain. Utility procurement was a real bottleneck. Bilateral contracts substantially improve project financial viability. The problem was real. The solution was incomplete.

Power purchase agreements resolved commercial risk for developers. They did not resolve interconnection queue positions, transmission availability, grid operator market rules, or generation siting constraints. Constellation's restart of Three Mile Island Unit 1—backed by a Microsoft power purchase agreement, a completed NRC restart approval, and strong commercial commitment—was delayed by transmission constraints that the commercial contract could not address.[13] The ability to contract for power did not produce the ability to deliver it.

Hyperscalers' response illustrates the logic of navigating a distributed decision system. Rather than persisting with the single-domain approach, they expanded engagement across the system. They began investing directly in generation developers, influencing project design and risk allocation from within the financial and commercial domains rather than only from the demand side. They engaged with grid operators and market designers on large-load tariff structures and interconnection processes, entering the technical and infrastructure domain. They submitted interconnection requests across multiple regions simultaneously, using the queue as a discovery mechanism for where capacity could actually be delivered on acceptable timelines.[14,15] Several have engaged in permitting reform advocacy, recognizing that regulatory throughput constraints govern project timelines they cannot control commercially.

This evolution was an adaptive response to constraints encountered sequentially as the commercial intervention revealed the domains that remained unresolved. The hyperscaler case demonstrates, in near-real-time, what the framework predicts: a well-capitalized actor entering through a single domain encounters that domain's limits, discovers that the constraints persist elsewhere, and progressively expands its engagement until it is operating across the full system—subject, at each step, to the same requirement for convergence across independent decision processes.

Their presence in energy markets has made the system's constraints more visible, more quickly. That is the correct characterization of what a new institutional actor does to a distributed decision system. It does not resolve it. It reveals it.

These four cases differ in detail but converge on the same structural outcome: interventions alter the location and timing of constraints rather than eliminating them. The system is more resilient to single-domain intervention than any domain-level analysis suggests, precisely because its constraints are distributed, interactive, and partially self-generating.

Conclusion

Large-scale infrastructure projects consistently fail in ways that domain-level analysis does not predict. Each approval may be obtained. Each stakeholder assessment may be sound. Each intervention may succeed within its own scope. And yet projects stall, collapse during execution, or fail to be delivered at the system level. The explanation is structural.

The framework developed in this report identifies that structure. Distributed decision authority means that no single actor determines outcomes and no single intervention is sufficient to change them. Non-fungible approvals mean that progress is conjunctive: every condition but one resolved advances nothing. Fragmented risk means that the same event is evaluated differently across domains, that evidence from one domain does not resolve concerns in another, and that risk mitigation in one domain frequently constitutes risk transfer to another. Binding feasibility constraints mean that authorization and execution depend on different conditions, some of which can become binding after approvals are secured.

These characteristics produce a specific, falsifiable condition: in systems with genuinely distributed decision authority, alignment is inherently transient, coordination is inherently fragile, and project realization is structurally uncertain even when alignment appears to have been achieved. This condition is managed rather than solved. It weakens where authority consolidates and intensifies where authority fragments. It does not disappear from any system in which independent stakeholders hold non-substitutable approvals.

The practical implications are direct. Risk allocation is a design choice with system-level consequences: risk transferred to actors without the capacity to absorb it re-emerges as instability in a different domain, at a different stage, often attributed to a different cause. Interventions that target a visible constraint without tracing how that constraint will reconstitute when displaced are incomplete—successful within their scope, inconsequential at the system level. And actors entering the system through a single domain will eventually encounter the full set of constraints, regardless of their financial scale or strategic sophistication. The question is whether they recognize the structure before or after significant resources have been committed on the wrong premise.

The framework does not supply a general solution. The condition it describes cannot be fully resolved while decision authority remains genuinely stratified. What the framework offers is a more accurate description of the problem, which is a prerequisite for designing interventions that can actually change it. Solutions designed for a system governed by a dominant constraint will not work in a system governed by joint satisfiability across independent domains. Understanding the difference is where effective analysis begins.

References

[1] Eliasson, Jonas & Fosgerau, Mogens, 2013. "Optimism bias in project appraisal: deception or selection?," Transportation Research B, vol 57, pages 105-113. DOI: 10.1016/j.trb.2013.09.005

[2] Nuclear Regulatory Commission (NRC), "Draft Environmental Impact Statement for Combined Licenses for Virgil C. Summer Nuclear Station Units 2 and 3," NUREG-1939 Vol. 1, Apr. 2010. [Online]. Available: https://www.regulations.gov/document/NRC-2008-0441-0031

[3] Heinz J. Mueller, "EPA Review and Comments Draft Environmental Impact Statement (DEIS) for the Combined Licenses (COLs) for Virgil C. Summer Nuclear Station (VCSNS) Units 2 and 3," July 10, 2010. [Online]. Available: https://www.regulations.gov/comment/NRC-2008-0441-0048

[4] U.S. Department of Energy, "About Community Benefits Plans," Energy.gov. https://www.energy.gov/infrastructure/about-community-benefits-plans [Retrieved 07/01/2026]

[5] U.S. Department of Energy, "Community Benefits Plan Frequently Asked Questions (FAQs)," Energy.gov. https://www.energy.gov/clean-energy-infrastructure/community-benefits-plan-frequently-asked-questions-faqs

[6] U.S. Department of Energy, "Justice40 Initiative," Energy.gov. https://www.energy.gov/justice/justice40-initiative [Retrieved 07/01/2026]

[7] Neutron Bytes, "Diablo Canyon Agrees to Coastal Commission Land Use Plan," Neutron Bytes, Dec. 14, 2025. https://neutronbytes.com/2025/12/14/diablo-canyon-agrees-to-coastal-commission-land-use-plan/

[8] WINDExchange, "Wind Energy Community Benefits Guide," WINDExchange. https://windexchange.energy.gov [Retrieved 07/01/2026]

[9] Sabin Center for Climate Change Law, "Community Benefits Agreements Database," Columbia Climate School. https://climate.law.columbia.edu. [Retrieved 07/01/2026]

[10] Alex Trembath, Elizabeth McCarthy, Lauren Teixeira, Alex Smith, and Marc Levitt, "How NEPA Litigation Obstructs Critical Projects," Breakthrough Institute, July 2025. [Online]. Available: https://thebreakthrough.org/issues/energy/the-procedural-hangover

[11] Google and Kairos Power, "New Nuclear Clean Energy Agreement with Kairos Power," Oct. 14, 2024.

[12] Microsoft and Constellation Energy, "Microsoft and Constellation Announce Agreement to Restart Three Mile Island Unit 1 / Crane Clean Energy Center," 2024.

[13] L. Kearney, "Delayed transmission projects blocking speedy Three Mile Island restart, Constellation says," Reuters, Apr. 06, 2026. [Online]. Available: https://www.reuters.com/business/energy/delayed-transmission-projects-blocking-speedy-three-mile-island-restart-2026-04-06/

[14] Amazon Web Services, "Amazon Signs Agreements to Support Development of Nuclear Energy Projects," 2024.

[15] Google, "How we're working with utilities to create a new model for clean energy," June 11, 2024. https://blog.google/company-news/outreach-and-initiatives/sustainability/google-clean-energy-partnership/ (accessed June 03, 2026).

THE BREAKTHROUGH INSTITUTE

BERKELEY, CA 94704

WWW.THEBREAKTHROUGH.ORG

X:@TheBTI